Layout-Corrector: Alleviating Layout Sticking Phenomenon in Discrete Diffusion Model

Shoma Iwai¹*[®], Atsuki Osanai²[®], Shunsuke Kitada²[®], and Shinichiro Omachi¹[®]

¹ Graduate School of Engineering, Tohoku University, Japan shoma.iwai.s4@dc.tohoku.ac.jp, shinichiro.omachi.b5@tohoku.ac.jp
² LY Corporation, Japan {atsuki.osanai, s.kitada}@lycorp.co.jp

Abstract. Layout generation is a task to synthesize a harmonious layout with elements characterized by attributes such as category, position, and size. Human designers experiment with the placement and modification of elements to create aesthetic layouts, however, we observed that current discrete diffusion models (DDMs) struggle to correct inharmonious layouts after they have been generated. In this paper, we first provide novel insights into layout sticking phenomenon in DDMs and then propose a simple yet effective layout-assessment module Layout-Corrector, which works in conjunction with existing DDMs to address the layout sticking problem. We present a learning-based module capable of identifying inharmonious elements within layouts, considering overall layout harmony characterized by complex composition. During the generation process, Layout-Corrector evaluates the correctness of each token in the generated layout, reinitializing those with low scores to the ungenerated state. The DDM then uses the high-scored tokens as clues to regenerate the harmonized tokens. Layout-Corrector, tested on common benchmarks, consistently boosts layout-generation performance when in conjunction with various state-of-the-art DDMs. Furthermore, our extensive analysis demonstrates that the Layout-Corrector (1) successfully identifies erroneous tokens, (2) facilitates control over the fidelitydiversity trade-off, and (3) significantly mitigates the performance drop associated with fast sampling.

Keywords: Layout Generation · Discrete Diffusion Model

1 Introduction

Creating a layout is one of the most crucial tasks involving human labor when designing [1], and there are a wide variety of applications, including academic papers [50], application user interfaces [8], and advertisements [45]. Layout generation has been formulated as a task that determines a set of elements that consist of categories, positions, and sizes [33, 40]. In response, layout-generation methods on deep learning have shown remarkable performance, and in particular, discrete generative models such as masked language modeling [9]-based

 $^{^{\}star}\,$ This work was done during the first author's internship at LY Corporation.

methods [5, 26, 43] and discrete diffusion models (DDMs) [13, 20, 21, 47] are the current state-of-the-art (SoTA).

To create an aesthetically pleasing layout, human designers typically modify layouts through trial and error. However, we found that even SoTA DDMs can not update elements in layouts once they have been generated, *i.e.*, *layout sticking*. An intuitive example of the sticking behavior is depicted in Fig. 1, where inharmonious elements that arose during generation persist until the final generated result. While former studies [6,24,35,36] tried to refine these elements in the post-processing phase that minimizes the rule-based costs such as alignment, they could not capture higher-order senses that determine layout aesthetics.

In the image generation domain, non-autoregressive (Non-AR) decoding methods with an external critic have demonstrated remarkable performance [29, 30]. The module identifies deviated visual tokens from the real distribution and reset them to resample. Reviewing the success of the masked image modeling [16], a few visual clues can provide plenty of information to identify erroneous tokens. On the other hand, the layout domain has different characteristics; (i) unlike images with a fixed and enough number of tokens (e.g., 16×16 patches), the number of layout elements is small and varies across samples (e.g., 1 to)25 elements), and (ii) as the element composed of multiple attributes, then



Fig. 1: Intuitive overview of Layout-Corrector. Conventional generative models cannot modify the elements once they have been generated. Layout-Corrector works in conjunction with DDMs to identify inharmonious elements in the generative process and initialize them to enhance regeneration towards a harmonized layout.

partially observed tokens do not provide enough clues. Thus, it is non-trivial whether the technique in the vision can apply to the layout generation.

In this paper, we propose a simple yet effective approach, named *Layout-Corrector*, to address the layout sticking problem. It works as an external module that evaluates each token's correctness score in a learning-based manner, aiming to identify erroneous tokens in a layout. As shown in Fig. 1, during the generation process, tokens with low correctness scores are reset to the ungenerated state (*i.e.*, [MASK]). Then, a DDM regenerates them using the remaining high-scored tokens as clues. Additionally, to deal with the characteristics of the layout tokens mentioned above, we propose a new objective and application schedule that accommodates variable numbers of elements while providing reliable layout cues.

We conducted extensive experiments on Layout-Corrector using three benchmarks [8, 45, 50]. When in conjunction with strong baselines [5, 13, 21], Layout-Corrector significantly enhanced their performance in both unconditional and conditional generation tasks. Both quantitative and qualitative evaluations confirmed that our approach effectively corrects inharmonious layout elements, addressing the challenge present in SoTA DDMs. By adjusting the application schedule of the corrector, we also achieved enhanced control over the fidelitydiversity and speed-quality trade-offs, demonstrating Layout-Corrector's versatility across different application scenarios.

Our contributions are summarized as follows: (1) We empirically demonstrate that current SoTA DDMs struggle to correct inharmonious elements in layouts; however, they can effectively correct them when erroneous elements are initialized to the ungenerated state, [MASK]. (2) We propose Layout-Corrector for evaluating the correctness score of each element and resetting the element with a lower score to [MASK], enabling DDMs to regenerate improved layouts. (3) We confirm consistent improvements by applying Layout-Corrector to various DDMs. We also analyze the behavior of Layout-Corrector and demonstrate that it enhances fidelity-diversity and speed-quality trade-offs.

2 Related Works

Layout Generation. Automatic layout generation [33, 40] is a task involving the assignment of positions and categories to multiple elements, which has diverse applications in design areas like application user interfaces and academic papers [8, 14, 19, 45, 46, 48-51]. This task includes unconditional and conditional generation, considering user constraints, *e.g.*, partially specified elements.

Early layout generation research explored classical optimization [36, 37] and generative models such as GAN [12]-based models [31,49,51] and VAE [10]-based models [23,45]. Following the success in NLP, Transformer-based approaches [44] were proposed. Auto-regressive (AR) models [2,15] iteratively generate layouts, however, struggle with conditional generation [26]. Non-AR models [5, 26, 43] overcome this difficulty by using a bidirectional architecture, where user-defined conditions serve as clues to complete blank tokens. Recently, diffusion modelbased [18,41] layout generation methods in both continuous [4, 28] and discrete spaces [20,21,47] have been developed. To enable unconditional and conditional generation within a single framework, it is essential for models to process both discrete and continuous data present in elements. DDMs can accommodate both data types by quantizing geometric attributes into a binned space.

Discrete Diffusion Models. D3PM [3] introduces the special token [MASK], where regular tokens are absorbed into [MASK] through a forward process. Based on this, Non-AR models, such as MaskGIT [5], can be understood as a subclass of DDMs. MaskGIT introduces a scheduled masking rate, akin to the diffusion process in D3PM. It also adopts the parallel decoding [11] based on the token confidence, serving as a deterministic denoising process. To address the issue of non-regrettable decoding strategy, DPC and Token-Critic [29, 30] introduce an external module to mitigate discrepancies between training and inference distributions. VQDiffusion [13] facilitates transitions between regular tokens in addition to [MASK], while LayoutDM [21] advances the diffusion process to allow transitions within the same modality. LayoutDiffusion [47] introduces a mild corruption process that considers the continuity of geometric attributes. For

the layout generation, we explore the potential of the correction during the generation process to alleviate layout token sticking problem.

Layout Correction. There are several studies aimed at layout modification. In optimization-based methods [24, 35, 36], layouts are refined to minimize hand-crafted costs, such as alignment score. RUITE [38] and LayoutFormer++ [22] learn to restore the original layout from noisy input. LayoutDM [21] proposes a logit adjustment under the constraints of noisy layouts. While previous research has focused on layout refinement, our method aims to correct layouts during the generation process. Compared to rule-based optimization, our approach achieves superior performance while preserving the distribution of the generated results. Please refer to the supplementary material for details.

3 Method

In Sec. 3.1, we first provide a brief overview of DDMs for layout-generation [21] and examine the potential for layout correction in Sec. 3.2. We then present Layout-Corrector in Sec. 3.3 and explain its application across diverse layout-generation tasks in Sec. 3.4.

3.1 Layout Generation Models

A layout is represented as a set of elements, where an element consists of category, position, and size. Following previous studies [2, 15, 26], we use the quantized expression for geometric attributes. Defining $l_i = (c_i, x_i, y_i, w_i, h_i)$, a layout l with $N \in \mathbb{N}$ elements is expressed as $l_N = (l_1, \dots, l_N)$, where $c_i \in \{1, \dots, C\}$ denotes the category (e.g., text, button), $(x_i, y_i, w_i, h_i) \in \{1, \dots, B\}^4$ represents the center position and size of *i*-th element, and $B \in \mathbb{N}$ denotes the number of bins. Under this representation, we review DDMs to gain insight into the behavior of the generation process, as discussed in Sec. 3.2.

Let T represent the total number of timesteps in the corruption process. We consider a scalar variable z_t with $K \in \mathbb{N}$ classes at t, where $z_t \in \{1, \ldots, K\}$. Here, z_t substitutes an attribute of an element. Following LayoutDM [21], we include the special tokens [PAD] and [MASK], resulting in (K + 2) classes. Here, [PAD] token is employed to fill the empty element, achieving variable length generation. [MASK] token denotes the absorbing state, to which tokens converge through the diffusion process. Using a transition matrix $\mathbf{Q}_t \in [0, 1]^{(K+2) \times (K+2)}$, we can define a transition probability from z_{t-1} to z_t as follows:

$$q(z_t|z_{t-1}) = \boldsymbol{v}(z_t)^{\top} \mathbf{Q}_t \boldsymbol{v}(z_{t-1}), \qquad (1)$$

where $\boldsymbol{v}(z_t) \in \{0,1\}^{K+2}$ is a one-hot vector of z_t . Due to the Markov property, a transition from z_0 to z_t is similarly written as: $q(z_t|z_0) = \boldsymbol{v}(z_t)^{\top} \bar{\mathbf{Q}}_t \boldsymbol{v}(z_0)$, where $\bar{\mathbf{Q}}_t = \mathbf{Q}_t \mathbf{Q}_{t-1} \cdots \mathbf{Q}_1$. Applying the Markov property $q(z_{t-1}|z_t, z_0) = q(z_{t-1}|z_t)$, we can obtain the posterior distribution $q(z_{t-1}|z_t, z_0)$ (Eq. (5) in [13]).

For the reverse process, we compute a conditional distribution $p_{\theta}(\mathbf{z}_{t-1}|\mathbf{z}_t) \in [0,1]^{N \times (K+2)}$. Categorical variable \mathbf{z}_{t-1} is sampled from this distribution. As



(a) Top: $\bar{\beta}_t$ scheduling for timestep t. Bottom: (h plot of token-sticking-rate (TSR), representing T degree of token matching between z_0 and z_t .

(b) Success rate on token-correction task for Token-replace and Mask-replace strategies. Color of bar corresponds to legend in Fig. 2a.

Fig. 2: Results of preliminary experiments on Rico test set [8]. (a) While $\bar{\beta}_{t,K} = (K+1)\bar{\beta}_t = \epsilon \ (\ll 1)$ is affected by token-sticking, $\bar{\beta}_{t,K} > \epsilon$ alleviates it. (b) The results indicate that LayoutDM can restore the original tokens from [MASK]; however, recovery from regular tokens proves challenging. Please refer to Supp. for further results.

proposed in a previous study [3], we use the re-parametrization trick and obtain a posterior distribution as $p_{\theta}(\boldsymbol{z}_{t-1}|\boldsymbol{z}_t) \propto \sum_{\tilde{\boldsymbol{z}}_0} q(\boldsymbol{z}_{t-1}|\boldsymbol{z}_t, \tilde{\boldsymbol{z}}_0) \ \tilde{p}_{\theta}(\tilde{\boldsymbol{z}}_0|\boldsymbol{z}_t)$, where $\tilde{p}_{\theta}(\tilde{\boldsymbol{z}}_0|\boldsymbol{z}_t)$ is a neural network that predicts the noiseless token distribution at t = 0. Following previous studies [3,13,21], we employ the hybrid loss of variational lower bound and auxiliary denoising loss.

The design of transition matrix \mathbf{Q}_t is pivotal in defining the corruption process. The token transition is categorized into three types: (1) keeping the current token, (2) replacing the token with other tokens, and (3) replacing the token with [MASK]. For each, we employ the probabilities $(\alpha_t, \beta_t, \gamma_t)$. Hence, using $\mathbf{Q}'_t = \alpha_t \mathbb{I} + \beta_t \mathbf{1} \mathbf{1}^\top \in [0, 1]^{(K+1) \times (K+1)}$, $\mathbf{Q}_t \in [0, 1]^{(K+2) \times (K+2)}$ is defined as:

$$\mathbf{Q}_t = \begin{bmatrix} \mathbf{Q'}_t & \mathbf{0} \\ \gamma_t \cdots \gamma_t & 1 \end{bmatrix}.$$
(2)

The cumulative transition matrix \mathbf{Q}_t can be computed in the closed form as $\bar{\mathbf{Q}}_t \boldsymbol{v}(x_0) = \bar{\alpha}_t \boldsymbol{v}(x_0) + (\bar{\gamma}_t - \bar{\beta}_t) \boldsymbol{v}(K+2) + \bar{\beta}_t$, where $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, $\bar{\gamma}_t = 1 - \prod_{i=1}^t (1 - \gamma_i)$, $\bar{\beta}_{t,K} = (K+1)\bar{\beta}_t = 1 - \bar{\alpha}_t - \bar{\gamma}_t$, and $\boldsymbol{v}(K+2)$ denotes the one-hot representation of [MASK]. A transition with $\bar{\beta}_t > 0$ introduces a layout inconsistency. Since the corresponding DDM is trained to correct such mismatches, we expect that it can update erroneous tokens caused in the generation process. However, in the following section, we will demonstrate that the scheduling of $\bar{\beta}_t$ is suboptimal for layout correction.

3.2 Preliminary: Potential of Token Correction in DDM

We explore token correction with DDMs, specifically focusing on LayoutDM [21], by assessing the impact of the $\bar{\beta}_t$ schedule. For $\epsilon \ll 1$ and $\bar{\beta}_{t,K} = \epsilon$ for any t, $p_{\theta}(\boldsymbol{z}_{t-1}|\boldsymbol{z}_t)$ struggles to correct tokens, due to the diffusion process not facilitating token replacement, except for [MASK]. This limitation is analogous to those seen with parallel decoding methods used in MaskGIT [5]. A possible solution is to increase $\bar{\beta}_{t,K} > \epsilon$ to promote transition between regular tokens. To verify the



Fig. 3: The details of Layout-Corrector. Top: training procedure of Layout-Corrector, where the pre-trained DDM is fixed. Bottom: sampling process with Layout-Corrector. We execute the generation and correction process in the purple box iteratively.

effect of $\beta_{t,K}$, we compare the token-sticking-rate (TSR) in the reverse process, which measures the proportion of tokens at z_0 that remain unchanged from z_t . As depicted in Fig. 2a, $\bar{\beta}_{t,K} = \epsilon$ leads to TSR $\simeq 100\%$ at most t, indicating token sticking. In contrast, when $\bar{\beta}_{t,K} > \epsilon$, the TSR is reduced below 100%, indicating that $p_{\theta}(z_{t-1}|z_t)$ can update tokens during generation process.

We next evaluate the DDM's error-correction capability by simulating replacements of three randomly selected tokens in a sequence with either [MASK] or other tokens, and then observing the model's ability to restore the original tokens. These methods are referred to as *Mask-replace* and *Token-replace*, respectively. In this setup, LayoutDM executes the reverse step from timestep t = 10 to 1. Our metric is the success rate of token recovery, which is deemed successful if recovery is complete, and we assess this across different $\bar{\beta}_t$ schedules. Fig. 2b demonstrates that Token-replace with $\bar{\beta}_{t,K} > \epsilon$ is moderately more effective than when $\bar{\beta}_{t,K} = \epsilon$. However, Mask-replace exhibits significant improvements over Token-replace. This finding motivated us to develop Layout-Corrector, which resets inharmonious tokens to [MASK].

3.3 Layout-Corrector

To enhance the replacement of the erroneous tokens with [MASK], we introduce Layout-Corrector. Functioning as a quality assessor, Layout-Corrector evaluates the correctness of each token in a layout during the generation process. The tokens with lower correctness scores are replaced with [MASK], and the updated tokens are fed back into the DDM. Therefore, Layout-Corrector can explicitly prompt the DDM to modify the erroneous tokens.

Architecture. Evaluating the correctness of each token requires Layout-Corrector to consider the relationships between elements in a layout. To this end, we use a Transformer [44] encoder to capture global contexts, as shown in Fig. 3. We first apply a multi-layer perceptron (MLP) to fuse five tokens (c, x, y, w, h) of each element, obtaining N element embeddings. These embeddings are processed by a transformer encoder, producing five-channel outputs. Each channel corresponds

to the correctness score of each token in the element. Since the layout elements are order-agnostic, we eliminate positional encoding to avoid unintended biases. **Training.** The objective of Layout-Corrector is to detect erroneous tokens during the generation process. To achieve this, we train Layout-Corrector as a binary classifier with a pre-trained DDM, which is frozen during training, as shown in Fig. 3. Given an original layout z_0 and t, a forward process is applied to obtain a distribution $q(z_t|z_0)$, and DDM estimates the distribution $p_{\theta}(z_{t-1}|z_t)$. Then, for a [MASK]-free token sequence \hat{z}_{t-1} sampled from $p_{\theta}(z_{t-1}|z_t)$, Layout-Corrector evaluates the correctness score $p_{\phi}(\hat{z}_{t-1}, t) \in [0, 1]^{5N}$ for each token in \hat{z}_{t-1} that are originally masked in z_t , we train Layout-Corrector to predict whether each token in \hat{z}_{t-1} aligns with the corresponding original token in z_0 , as in [7]. It encourages the Layout-Corrector to evaluate the correctness of each token directly. Specifically, we use binary cross-entropy (BCE) loss:

$$\mathcal{L}_{\text{Corrector}} = \text{BCE}(\boldsymbol{m}, p_{\phi}(\hat{\boldsymbol{z}}_{t-1}, t)), \qquad (3)$$

where $m^{(i)} = 1$ if $\hat{z}_{t-1}^{(i)} = z_0^{(i)}$, otherwise $m^{(i)} = 0$. Through the training, Layout-Corrector learns to identify erroneous tokens that disturb the layout harmony.

3.4 Generating Layout with the Layout-Corrector

Unconditional Generation. As shown in Fig. 3, all tokens z_T are initialized with [MASK], and the final output is obtained at t = 0. At timestep t, a DDM predicts the distribution $p_{\theta}(z_{t-1}|z_t)$, from which we sample a [MASK]-free token sequence \hat{z}_{t-1} . Layout-Corrector then assesses the correctness scores $p_{\phi}(\hat{z}_{t-1}, t)$. We add Gumbel noise to $p_{\phi}(\hat{z}_{t-1}, t)$ to introduce randomness into the token selection. Then, we mask tokens whose scores are lower than a threshold θ_{th} . Another possible way is to choose tokens with the lowest $5N \cdot \bar{\gamma}_t$ scores, similar to [29, 30], where $\bar{\gamma}_t$ is the mask ratio at t (see Sec. 3.1). However, it may mask high-quality tokens when the majority have higher scores, leading to diminished cues for the DDM. The threshold mitigates this issue by selectively masking only those tokens with lower scores, thus preserving reliable cues for regeneration.

Conditional Generation. Layout-Corrector is versatile and can be seamlessly used for various conditional generation tasks without specialized training or finetuning. Given a set of partially known tokens, *e.g.*, element categories or sizes, the goal of conditional generation is to estimate the remaining unknown tokens. Following LayoutDM [21], we utilize the known condition tokens as an initial state for the generation process and maintain these tokens at each *t*. When Layout-Corrector is applied, a correctness score of 1 is assigned to the conditional tokens. In this way, Layout-Corrector encourages the DDM to modify erroneous tokens while ensuring that the known tokens are preserved.

Corrector Scheduling. Layout-Corrector can be applied at any t during the generation process. Unlike existing methods [29, 30], which apply the external assessor at every t, we selectively apply Layout-Corrector at specific timesteps. Remarkably, alongside LayoutDM [21], Layout-Corrector enhances generation

Table 1: Performance comparison of baseline models with/without external assessor on unconditional generation. *Arch.* represents the architecture of the discrete generative model. Metrics improved by the external module are highlighted in **bold**.

		Rico [8]		Crello [45]			PubLayNet [50]			
Model	Arch.	FID↓	Precision [↑]	Recall↑	FID↓	Precision [↑]	Recall↑	FID↓	Precision↑	Recall↑
MaskGIT [5]	Non-AR	70.37	0.793	0.437	35.32	0.802	0.376	34.23	0.587	0.460
+ Token-Critic [29]		15.65	0.682	0.843	7.59	0.735	0.815	17.55	0.579	0.825
+ Corrector (ours)		14.40	0.814	0.744	11.17	0.839	0.696	13.74	0.501	0.883
VQDiffusion [13]	DDMs	7.83	0.716	0.907	5.57	0.740	0.884	12.38	0.567	0.925
+ Token-Critic [29]		15.22	0.842	0.731	10.05	0.834	0.657	17.53	0.812	0.628
+ Corrector (ours)		5.29	0.809	0.898	4.70	0.793	0.842	9.89	0.699	0.903
LayoutDM [21]	DDMs	6.37	0.759	0.906	5.28	0.768	0.875	13.72	0.557	0.919
+ Token-Critic [29]		17.97	0.884	0.670	9.01	0.844	0.678	22.27	0.836	0.582
+ Corrector (ours)		4.79	0.811	0.891	4.36	0.822	0.851	11.85	0.711	0.890

quality with just three applications, effectively reducing additional forward operations during inference. Moreover, by adjusting the schedule of Layout-Corrector, we can modulate the fidelity-diversity trade-off of the generated layouts. Specifically, more frequent corrector applications enhance fidelity by removing a larger number of inharmonious tokens, while a more sparse schedule improves diversity. The experimental section provides a more detailed analysis on the schedules.

4 Experiments

4.1 Experimental Settings

Datasets. We evaluated Layout-Corrector on the following three challenging layout datasets over different domains: *Rico* [8] contains user interface designs for mobile applications. It contains 25 element categories such as text, button, and icon. *Crello* [45] consists of design templates for various formats, such as social media posts and banner ads. *PubLayNet* [50] comprises academic papers with 5 categories, such as table, image, and text. We follow the dataset splits presented in a previous study [21] for Rico and PubLayNet and use the official splits for Crello. We excluded layouts with more than 25 elements as in [21].

Evaluation Metrics. We used the following evaluation metrics: *Fréchet Inception Distance (FID)* [17] evaluates the similarity between distributions of generated and real data in the feature space using the feature extractor [24]. *Alignment (Align.)* [32] measures the alignment of elements in generated layouts. This metric is normalized by the number of elements, as in [24]. *Maximum IoU (Max-IoU)* [24] evaluates the similarity of the elements in bounding boxes of the same category, comparing the generated layouts to the ground truth. For fidelity and diversity, we used *Precision* and *Recall* [27].

Layout-generation Tasks. We evaluated our Layout-Corrector across three tasks: Unconditional generates a layout without any constraints. Category \rightarrow size + position ($C \rightarrow S + P$) generates a layout given only the category of each



Fig. 4: Accuracy of detecting erroneous tokens when three tokens are replaced randomly.



Fig. 5: Correctness scores against an extent of layout disruption controlled by various maximum transition steps.

element. Category + size \rightarrow position (C + S \rightarrow P) generates a layout given the category and size of each element.

Implementation Details. We used DDMs, *i.e.*, LayoutDM [21] and VQDiffusion [13], as well as a non-AR model *i.e.*, MaskGIT [5], as baseline models, and applied Layout-Corrector to them. Since non-ARs can be understood as a subclass of DDMs as discussed in Sec. 2, we can seamlessly apply Layout-Corrector to them. We used the publicly available pre-trained LayoutDM on Rico and PubLayNet, while we trained other models using LayoutDM implementation. Unless otherwise specified, the total timesteps T for LayoutDM and VQDiffusion were set to 100, and Layout-Corrector was applied at $t = \{10, 20, 30\}$, leading to a total of 103 forward operations. In MaskGIT [5], T = 10 and Layout-Corrector was applied at every t. For the threshold θ_{th} , we set it to 0.7 for LayoutDM and VQDiffusion, and 0.3 for MaskGIT. To train Layout-Corrector, we used AdamW [25, 34] with an initial learning rate of 5.0×10^{-4} and $(\beta_1, \beta_2) = (0.9, 0.98)$. Refer to supplementary material for more details.

4.2 Effectiveness of Layout-Corrector

To evaluate the applicability and effectiveness of Layout-Corrector with various discrete generative models, we applied our corrector and Token-Critic [29] to MaskGIT [5], VQDiffusion [13], and LayoutDM [21]. The results in Tab. 1 show that Layout-Corrector consistently improved FID across all tested models, confirming its effectiveness. In contrast, while Token-Critic [29] enhanced FID when applied to MaskGIT, its application to VQDiffusion and LayoutDM resulted in diminished performance. These results demonstrate that the direct application of Token-Critic can lead to suboptimal performance in layout generation, highlighting the importance of tailored approaches. Regarding fidelity and diversity, evaluated using Precision and Recall [27], we observed different trends between a non-AR model and DDMs. For MaskGIT, Layout-Corrector boosted both diversity and fidelity. The parallel decoding in MaskGIT keeps high-confidence tokens and rejects low-confidence ones. It often leads to stereotypical token patterns, resulting in high fidelity but low diversity. Layout-Corrector resets such patterns



conditional generation on different corrector schedules.

Fig. 6: FID-Precision trade-off on un- Fig. 7: Histogram of the width of bounding boxes on the Rico dataset on different corrector schedules.

while considering the overall harmony, thereby improving diversity without sacrificing fidelity. In the case of VQDiffusion and LayoutDM, Layout-Corrector increased fidelity while maintaining diversity. While the stochastic nature of DDMs promotes diversity, it can produce low-quality tokens. Layout-Corrector mitigates this issue by resetting these tokens, thereby enhancing fidelity.

4.3Analysis

Intrinsic Evaluation of Corrector. We assess Layout-Corrector's ability to detect erroneous tokens. To this end, we randomly replace three tokens in the ground truth with alternate ones using the test set of Rico dataset. The goal is for the corrector to identify these altered tokens. For Layout-Corrector and Token-Critic trained with the same LayoutDM, we evaluate the detection accuracy of these altered tokens, selecting the three tokens with the lowest corrector scores for comparison. Fig. 4 shows that Layout-Corrector outperforms Token-Critic due to the objective that directly estimates the correctness of tokens, underscoring its effectiveness in layout assessment.

Furthermore, we analyze the correlation between the degree of layout corruption and the correctness scores. To modulate the extent of disruption, we limit the maximum transition step for the geometric attributes when replacing. Fig. 5 depicts the average correctness scores for the three replaced and the other clean tokens within the corrupted layouts. We observe that a greater deviation from the original token leads to a lower correctness score against clean tokens. These results suggest that the corrector can measure the degree of discrepancy between the ideal and the actual layouts.

Corrector Scheduling: Impact on Fidelity and Diversity. We applied Layout-Corrector to LayoutDM with various schedules. Specifically, it is applied at $t = \{10\}, \{10, 20\}, \dots, \{10, 20, \dots, 90\}$, yielding nine distinct schedules. Fig. 6 shows the FID-Precision trade-offs with and without Layout-Corrector. It can be observed that varying the schedule effectively adjusts the FID-Precision trade-off. More frequent corrector applications enhance fidelity (*i.e.*, Precision), while decreasing diversity (*i.e.*, FID). These results align with Layout-Corrector's



Fig. 8: Speed-quality (FID) trade-off on unconditional generation. The Layout-Corrector mitigates the quality decline at smaller total timesteps.



Fig. 9: Average correctness scores on intermediate generated layouts of LayoutDM at each t across different total timesteps T' on Rico dataset.

role of resetting poor-quality tokens. Its frequent use leads to high-fidelity outputs; however, this also reduces the stochastic nature of DDMs, thus diminishing diversity. Moreover, it shows that our default schedule $t = \{10, 20, 30\}$ yields preferable FID, confirming the effectiveness of our schedule.

To further explore the impact of the corrector schedule, we analyzed the distribution of tokens' width attribute on Rico dataset. Fig. 7 compares the histograms of real and generated tokens under various corrector schedules. It reveals that more frequent correction amplifies the frequency trends of the original distribution. Concretely, values that are already common in the real data (e.q., w = 1.0 in Fig. 7) become more prevalent, while the occurrence of rarer values further diminishes. This observation aligns with the trend in Fig. 6, where frequent correction leads to higher fidelity but at the cost of reduced diversity. **Speed-Quality Trade-off.** Since generation speed is crucial in practical applications, we examined the speed-quality trade-off. To adjust the runtime of LayoutDM, we used the fast-sampling technique [3], which uses modulated distribution $p_{\theta}(\boldsymbol{z}_{t-\Delta}|\boldsymbol{z}_t)$ instead of $p_{\theta}(\boldsymbol{z}_{t-1}|\boldsymbol{z}_t)$. Δ is a step size, and the total steps are reduced to $T' = T/\Delta$. Regardless of Δ , Layout-Corrector is applied at $t = \{10, 20, 30\}$. Fig. 8 presents the results on $T' = \{20, 30, 50, 75, 100\}$. Compared with LayoutDM alone, Layout-Corrector improves the FID score with only a minimal increase in runtime, offering a superior trade-off. While the original LavoutDM's performance significantly degrades with a smaller T'. Lavout-Corrector effectively mitigates this issue by rectifying the misgenerated tokens, demonstrating the robustness in smaller T'. Notably, Layout-Corrector attains a competitive FID to that of the original LayoutDM (T' = 100) with just T' = 20.

To further analyze the benefits of Layout-Corrector for smaller T', we examined the correctness scores $p_{\phi}(\hat{z}_{t-1}, t)$ across different T'. Fig. 9 presents the average scores for intermediate tokens \hat{z}_t at each t on $T' = \{20, 50, 100\}$. We include the scores of real layouts for reference. Note that $p_{\phi}(\hat{z}_{t-1}, t)$ is affected by the input t because of the training procedure. For example, at smaller t, the corrupted tokens z_t become closer to z_0 ; therefore, most tokens in \hat{z}_{t-1} align

Table 2: Comparison with SoTA models on unconditional task. For $Align. \rightarrow$, values that more closely match those of real data are preferred, and we scale the values $100 \times$ for visibility. Best and second-best results are in **bold** and with underline, respectively.

		Rico [8]		Crello [45]		PubLayNet [50]	
Task	Model	FID↓	$\mathrm{Align.}{\rightarrow}$	$\mathrm{FID}{\downarrow}$	$\mathrm{Align.}{\rightarrow}$	$\mathrm{FID.}{\downarrow}$	$\mathrm{Align.}{\rightarrow}$
	DLT [28]	6.20	0.386	4.71	0.484	7.87	0.121
	LayoutTransformer [15]	7.63	0.068	5.93	0.305	13.90	0.127
7	LayoutDM [21]	6.39	0.223	5.28	0.279	13.69	0.185
ona	LayoutDM + Corrector	4.79	0.167	4.36	0.232	11.85	0.172
litio	Large models						
ŭ	LayoutDiffusion [47]	3.84	0.092	6.61	0.228	7.57	0.077
ICO	LayoutDM [*] [21]	4.93	0.146	4.40	0.315	10.92	0.158
Ur	$LayoutDM^* + Corrector$	4.23	0.127	$\overline{4.11}$	0.278	9.85	0.122
	Real data	1.85	0.109	2.32	0.338	6.25	0.0214



Fig. 10: Comparison of unconditional and conditional generation. In LayoutDM* + Corrector, Layout-Corrector is applied to the same intermediate states of LayoutDM* [21] at $t = \{10, 20, 30\}$ during generation process. Consequently, generation processes of both methods are identical from t = 100 to 30. While LayoutDM* generates unnatural elements, as highlighted in the blue dashed-line boxes, they are rectified by Layout-Corrector in LayoutDM* + Corrector. Refer to Supp. for more results.

with z_0 . Thus, the corrector learns to predict higher $p_{\phi}(\hat{z}_{t-1}, t)$ at smaller t. For the original LayoutDM (solid lines), there is a gap between scores at T' = 100and smaller T'. As indicated in Fig. 5, a lower $p_{\phi}(\hat{z}_{t-1}, t)$ suggests a greater discrepancy, resulting in inferior FID in Fig. 8. Conversely, by applying Layout-Corrector (dashed lines), low-scored tokens are reset, mitigating the gap in correctness scores at the smaller T'. It allows LayoutDM to leverage high-quality tokens as clues, effectively avoiding the deterioration in FID.

4.4 Comparison with State-of-the-Arts

Tabs. 2 and 3 present comparisons of Layout-Corrector with SoTA approaches for unconditional and conditional generation tasks, respectively. Following [21], we used FID and Alignment for the unconditional task, and Max IoU alongside

		Rico [8]		Crello [45]		PubLayNet [50]	
Task	Model	FID↓	Max-IoU↑	FID↓	Max-IoU↑	FID↓	Max-IoU↑
	LayoutGAN++ [24]	6.84	0.267	-	-	24.00	0.263
	DLT [28]	3.97	0.288	4.29	0.212	4.30	0.345
	LayoutTransformer [15]	5.57	0.223	6.42	0.203	14.10	0.272
•	LayoutDM [21]	3.51	0.276	4.04	0.197	7.94	0.309
$\mathrm{C}{\rightarrow}\mathrm{S+P}$	LayoutDM + Corrector	2.39	0.283	3.39	0.202	5.84	0.319
	Large models						
	LayoutNUWA [42]	2.52	0.445	-	-	6.58	0.385
	LayoutDiffusion [47]	1.13	0.357	4.68	0.253	3.09	0.351
	LayoutDM [*] [21]	2.12	0.302	3.04	0.206	6.25	0.322
	$LayoutDM^* + Corrector$	1.71	0.305	2.84	0.210	5.01	0.329
	LayoutGAN++ [24]	6.22	0.348	-	-	9.94	0.342
	DLT [28]	3.28	0.385	3.68	0.278	1.53	0.425
C+S→P	LayoutTransformer [15]	3.73	0.323	3.87	0.258	16.90	0.320
	LayoutDM [21]	2.17	0.390	3.55	0.248	4.22	0.380
	LayoutDM + Corrector	1.91	0.398	3.32	0.253	2.93	0.390
	Large models						
	LayoutNUWA [42]	2.87	0.564	-	-	3.70	0.483
	LayoutDM [*] [21]	1.29	0.460	2.71	0.269	2.69	0.408
	$LayoutDM^* + Corrector$	1.22	0.463	2.69	0.271	2.05	0.415

 Table 3: Comparison with SoTA models on conditional tasks. Best and second-best results are in **bold** and with <u>underline</u>, respectively.

FID for the conditional tasks. For the comparison, we chose the combination of LayoutDM and Layout-Corrector since it achieves the best performance in Tab. 1. For comparison with larger models [42, 47], we trained enlarged LayoutDM with 12 transformer layers, which is denoted as $LayoutDM^*$. Note that the architecture of the Layout-Corrector remains the same for $LayoutDM^*$.

As shown in Tabs. 2 and 3, our approach achieves superior or competitive FID scores compared with SoTA methods. While LayoutDiffusion [47] outperformed our approach on PubLayNet and Rico, Layout-Corrector achieves the best FID on Crello. On conditional tasks, although LayoutNUWA [42] achieved a notably higher Max IoU by using the substantially larger Code Llama 7B [39] and additional pre-training data, our method showed superior FID across the board. Overall, the consistent high performance of Layout-Corrector across a variety of tasks and datasets underscores its versatility and practical utility.

4.5 Qualitative Evaluation

Fig. 10 illustrates the layouts generated by different models under two tasks on the Rico and PubLayNet datasets. To demonstrate Layout-Corrector's impact, the results of LayoutDM* [21] and LayoutDM* + Corrector in the figure share the same intermediate states z_t in the generation process until the t when the corrector is first applied. While the overall structures of their outputs are similar, the enhancements from Layout-Corrector are clearly recognizable. For example, in the Rico dataset, Layout-Corrector successfully fixes the misalignments in LayoutDM's output. Moreover, the corrector rearranges overlapping elements in the PubLayNet. Compared with the other layout-generation methods, our approach consistently generates high-quality layouts, indicating its effectiveness.

4.6 Ablation Study

To validate each strategy in Layout-Corrector, we trained and tested the corrector with LayoutDM [21] under the following configurations. In w/o self-attention, the corrector uses only MLPs, lacking the ability to consider the global harmony of the layout. In Mask estimation, as in [29], the corrector predicts whether each token was originally

Table 4: Ablation study on Rico [8]with unconditional generation.

	T'	= 100	T' = 20			
	FID↓	$\overline{\text{Align.}} \rightarrow$	FID↓	$\text{Align.} \rightarrow$		
Layout-Corrector	4.79	0.167	6.84	0.159		
w/o self-attention	5.54	0.151	9.11	0.213		
Mask estimation	5.22	0.148	9.66	0.232		
Lowest-K	10.6	0.263	11.5	0.359		
Correct at every t	97.4	0.002	70.1	0.006		
Real Data	1.85	0.109	1.85	0.109		

masked rather than estimating if it aligns with the original one. In *Lowest-K*, we mask tokens with the lowest $5N \cdot \bar{\gamma}_t$ scores instead of using the threshold. In *Correct at every t*, the corrector is applied at every t during the reverse process. Except for this setting, the corrector was applied at $t = \{10, 20, 30\}$.

Tab. 4 shows the results in unconditional generation for $T' = \{100, 20\}$. Layout-Corrector achieved the best FID across different T'. In contrast, w/o selfattention resulted in inferior performance, showing the importance of capturing the relationship between elements. The results also demonstrate that the differences between Layout-Corrector and existing modules, *i.e.*, Token-Critic [29] and DPC [30], contribute to higher performance. As described in Sec. 3, the primary distinctions are (1) the training objective, (2) introducing threshold, and (3) selective scheduling. The results of *Mask estimation*, *Lowest-K*, and *Correcting* at every t indicate the effectiveness of each modification, respectively.

5 Conclusion

We introduced Layout-Corrector, a novel module working with a DDM-based layout-generation model. Our preliminary experiments highlighted (1) the tokensticking problem with DDMs and (2) the importance of masking mis-generated tokens to correct them. Based on these insights, we design Layout-Corrector to assess the correctness score of each token and replace the tokens with low scores with [MASK], guiding the generative model to correct these tokens. Our experiments showed that Layout-Corrector enhances the generation quality of various generative models on various tasks. Additionally, we have shown that it effectively controls the fidelity-diversity trade-off through its application schedule and mitigates the performance decline associated with fast sampling.

Limitations and Future Work. While Layout-Corrector adds marginal runtime, it increases memory usage and total parameter count. Our future work will aim to incorporate layout-specific mechanisms, e.g., element-relation embeddings [20], to improve Layout-Corrector's layout understanding capabilities.

Acknowledgment

We thank Dr. Kent Fujiwara and Dr. Hirokatsu Kataoka for their invaluable feedback and constructive suggestions on this paper.

References

- Agrawala, M., Li, W., Berthouzoz, F.: Design Principles for Visual Communication. Communications of the ACM 54(4), 60–69 (2011). https://doi.org/10.1145/ 1924421.1924439
- Arroyo, D.M., Postels, J., Tombari, F.: for Layout Generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13642-13652 (2021), https://openaccess.thecvf.com/content/CVPR2021/html/ Arroyo_Variational_Transformer_Networks_for_Layout_Generation_CVPR_ 2021_paper.html
- Austin, J., Johnson, D.D., Ho, J., Tarlow, D., Van Den Berg, R.: Structured Denoising Diffusion Models in Discrete State-Spaces. Advances in Neural Information Processing Systems 34, 17981–17993 (2021), https://proceedings.neurips.cc/ paper/2021/hash/958c530554f78bcd8e97125b70e6973d-Abstract.html
- 4. Chai, S., Zhuang, L., Yan, F.: LayoutDM: Transformer-based Diffusion Model for Layout Generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18349-18358 (2023), https://openaccess. thecvf.com/content/CVPR2023/html/Chai_LayoutDM_Transformer-Based_ Diffusion_Model_for_Layout_Generation_CVPR_2023_paper.html
- 5. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: MaskGIT: Masked Generative Image Transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11315-11325 (2022), https: //openaccess.thecvf.com/content/CVPR2022/html/Chang_MaskGIT_Masked_ Generative_Image_Transformer_CVPR_2022_paper.html
- Chen, J., Zhang, R., Zhou, Y., Chen, C.: Towards Aligned Layout Generation via Diffusion Model with Aesthetic Constraints. In: The Twelfth International Conference on Learning Representations (2024), https://openreview.net/forum?id= kJ0qp9Xdsh
- 7. Clark, K., Luong, M.T., Le, Q.V., Manning, C.D.: ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. In: International Conference on Learning Representations (2020), https://openreview.net/pdf?id= r1xMH1BtvB
- Deka, B., Huang, Z., Franzen, C., Hibschman, J., Afergan, D., Li, Y., Nichols, J., Kumar, R.: 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. pp. 845–854 (2017). https://doi.org/10.1145/3126594. 312665
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 4171–4186 (2019). https://doi.org/10.18653/v1/N19-1423
- Diederik P Kingma, M.W.: Auto-Encoding Variational Bayes. In: International Conference on Learning Representations (2014), https://openreview.net/forum? id=33X9fd2-9FyZd

- 16 Iwai et al.
- Ghazvininejad, M., Levy, O., Liu, Y., Zettlemoyer, L.: 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). pp. 6112–6121 (2019). https://doi.org/10.18653/v1/D19-1633
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative Adversarial Nets. Advances in neural information processing systems 27 (2014), https://papers.nips.cc/paper_files/ paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html
- Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector Quantized Diffusion Model for Text-to-Image Synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10696-10706 (2022), https://openaccess.thecvf.com/content/CVPR2022/html/ Gu_Vector_Quantized_Diffusion_Model_for_Text-to-Image_Synthesis_CVPR_ 2022_paper.html
- 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). https://doi.org/10.1145/3411764.3445117
- 15. Gupta, K., Lazarow, J., Achille, A., Davis, L.S., Mahadevan, V., Shrivastava, A.: LayoutTransformer: Layout Generation and Completion with Self-attention. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1004-1014 (2021), https://openaccess.thecvf.com/content/ICCV2021/html/ Gupta_LayoutTransformer_Layout_Generation_and_Completion_With_Self-Attention_ICCV_2021_paper.html
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked Autoencoders Are Scalable Vision Learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16000-16009 (June 2022), https://openaccess.thecvf.com/content/CVPR2022/html/He_Masked_ Autoencoders_Are_Scalable_Vision_Learners_CVPR_2022_paper.html
- 17. 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. Advances in neural information processing systems 30 (2017), https://papers. nips.cc/paper_files/paper/2017/hash/8a1d694707eb0fefe65871369074926d-Abstract.html
- Ho, J., Jain, A., Abbeel, P.: Denoising Diffusion Probabilistic Models. Advances in neural information processing systems 33, 6840– 6851 (2020), https://proceedings.neurips.cc/paper/2020/hash/ 4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html
- 19. Hsu, H.Y., He, X., Peng, Y., Kong, H., Zhang, Q.: PosterLayout: A New Benchmark and Approach for Content-aware Visual-Textual Presentation Layout. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6018-6026 (2023), https://openaccess.thecvf.com/content/CVPR2023/ html/Hsu_PosterLayout_A_New_Benchmark_and_Approach_for_Content-Aware_ Visual-Textual_Presentation_CVPR_2023_paper.html
- 20. Hui, M., Zhang, Z., Zhang, X., Xie, W., Wang, Y., Lu, Y.: Unifying Layout Generation with a Decoupled Diffusion Model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1942-1951 (2023), https://openaccess.thecvf.com/content/CVPR2023/html/Hui_Unifying_Layout_Generation_With_a_Decoupled_Diffusion_Model_CVPR_2023_paper.html

- Inoue, N., Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: LayoutDM: Discrete Diffusion Model for Controllable Layout Generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10167-10176 (2023), https://openaccess.thecvf.com/content/CVPR2023/html/Inoue_LayoutDM_Discrete_Diffusion_Model_for_Controllable_Layout_Generation_CVPR_2023_paper.html
- 22. Jiang, Z., Guo, J., Sun, S., Deng, H., Wu, Z., Mijovic, V., Yang, Z.J., Lou, J.G., Zhang, D.: 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. pp. 18403-18412 (2023), https://openaccess.thecvf.com/content/CVPR2023/html/Jiang_LayoutFormer_Conditional_Graphic_Layout_Generation_via_Constraint_Serialization_and_Decoding_CVPR_2023_paper.html
- 23. Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: LayoutVAE: Stochastic Scene Layout Generation From a Label Set. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9895-9904 (2019), https: //openaccess.thecvf.com/content_ICCV_2019/html/Jyothi_LayoutVAE_ Stochastic_Scene_Layout_Generation_From_a_Label_Set_ICCV_2019_paper. html
- Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: Constrained Graphic Layout Generation via Latent Optimization. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 88–96 (2021). https://doi.org/10.1145/ 3474085.347549
- Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: International Conference on Learning Representations (2015), https://openreview.net/ forum?id=8gmWwjFyLj
- Kong, X., Jiang, L., Chang, H., Zhang, H., Hao, Y., Gong, H., Essa, I.: BLT: Bidirectional Layout Transformer for Controllable Layout Generation. In: European Conference on Computer Vision. pp. 474–490 (2022). https://doi.org/10.1007/978-3-031-19790-1_29
- Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., Aila, T.: Improved Precision and Recall Metric for Assessing Generative Models. Advances in Neural Information Processing Systems 32 (2019), https://papers.nips.cc/paper_files/ paper/2019/hash/0234c510bc6d908b28c70ff313743079-Abstract.html
- 28. Levi, E., Brosh, E., Mykhailych, M., Perez, M.: DLT: Conditioned layout generation with Joint Discrete-Continuous Diffusion Layout Transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2106-2115 (2023), https://openaccess.thecvf.com/content/ICCV2023/html/Levi_ DLT_Conditioned_layout_generation_with_Joint_Discrete-Continuous_ Diffusion_Layout_Transformer_ICCV_2023_paper.html
- Lezama, J., Chang, H., Jiang, L., Essa, I.: Improved Masked Image Generation with Token-Critic. In: European Conference on Computer Vision. pp. 70–86 (2022). https://doi.org/10.1007/978-3-031-20050-2_5
- Lezama, J., Salimans, T., Jiang, L., Chang, H., Ho, J., Essa, I.: Discrete Predictor-Corrector Diffusion Models for Image Synthesis. In: The Eleventh International Conference on Learning Representations (2023). https://doi.org/https:// openreview.net/forum?id=VM8batVBWvg
- 31. Li, J., Yang, J., Hertzmann, A., Zhang, J., Xu, T.: LayoutGAN: Generating Graphic Layouts with Wireframe Discriminators. In: The Seventh International

Conference on Learning Representations (2019), https://openreview.net/forum? id=HJxB5sRcFQ

- 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). https://doi.org/10.1109/TVCG. 2020.299933
- 33. Lok, S., Feiner, S.: A Survey of Automated Layout Techniques for Information Presentations. Proceedings of SmartGraphics 2001, 61-68 (2001). https: //doi.org/https://hci.stanford.edu/courses/cs448b/papers/LokFeiner_ layoutsurvey.pdf
- Loshchilov, I., Hutter, F.: Decoupled Weight Decay Regularization. In: The Seventh International Conference on Learning Representations (2018), https:// openreview.net/forum?id=Bkg6RiCqY7
- 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). https://doi.org/10.1145/1964921.1964982
- 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). https://doi.org/10.1109/TVCG.2014.48
- O'Donovan, P., Agarwala, A., Hertzmann, A.: DesignScape: Design with Interactive Layout Suggestions. In: Proceedings of the 33rd annual ACM conference on human factors in computing systems. pp. 1221–1224 (2015). https://doi.org/ 10.1145/2702123.2702149
- Rahman, S., Sermuga Pandian, V.P., Jarke, M.: RUITE: Refining UI Layout Aesthetics Using Transformer Encoder. In: 26th International Conference on Intelligent User Interfaces-Companion. pp. 81–83 (2021). https://doi.org/10.1145/3397482.3450716
- Roziere, B., Gehring, J., Gloeckle, F., Sootla, S., Gat, I., Tan, X.E., Adi, Y., Liu, J., Remez, T., Rapin, J., et al.: Code Llama: Open Foundation Models for Code. arXiv preprint arXiv:2308.12950 (2023). https://doi.org/10.48550/arXiv.2308.12950
- 40. Shi, Y., Shang, M., Qi, Z.: Intelligent layout generation based on deep generative models: A comprehensive survey. Information Fusion p. 101940 (2023). https: //doi.org/10.1016/j.inffus.2023.101940
- 41. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In: International conference on machine learning. pp. 2256-2265 (2015), https://proceedings.mlr. press/v37/sohl-dickstein15.html
- 42. Tang, Z., Wu, C., Li, J., Duan, N.: LayoutNUWA: Revealing the Hidden Layout Expertise of Large Language Models. In: The Twelfth International Conference on Learning Representations (2024), https://openreview.net/forum?id= qCUWVT0Ayy
- Turgutlu, K., Sharma, S., Kumar, J.: LayoutBERT: Masked Language Layout Model for Object Insertion. arXiv preprint arXiv:2205.00347 (2022). https:// doi.org/10.48550/arXiv.2205.00347
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention Is All You Need. Advances in neural information processing systems **30** (2017), https://papers.nips.cc/paper_files/paper/ 2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html
- 45. Yamaguchi, K.: CanvasVAE: Learning To Generate Vector Graphic Documents. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp.

5481-5489 (2021), https://openaccess.thecvf.com/content/ICCV2021/html/ Yamaguchi_CanvasVAE_Learning_To_Generate_Vector_Graphic_Documents_ ICCV_2021_paper.html

- Yang, X., Mei, T., Xu, Y.Q., Rui, Y., Li, S.: Automatic Generation of Visual-Textual Presentation Layout. ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM) 12(2), 1–22 (2016). https://doi.org/ 10.1145/2818709
- 47. Zhang, J., Guo, J., Sun, S., Lou, J.G., Zhang, D.: LayoutDiffusion: Improving Graphic Layout Generation by Discrete Diffusion Probabilistic Models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7226-7236 (2023), https://openaccess.thecvf.com/content/ICCV2023/ html/Zhang_LayoutDiffusion_Improving_Graphic_Layout_Generation_by_ Discrete_Diffusion_Probabilistic_Models_ICCV_2023_paper.html
- Zhao, T., Chen, C., Liu, Y., Zhu, X.: guigan: Learning to Generate GUI Designs Using Generative Adversarial Networks. In: 2021 IEEE/ACM 43rd International Conference on Software Engineering. pp. 748–760 (2021). https://doi.org/10. 1109/ICSE43902.2021.00074
- Zheng, X., Qiao, X., Cao, Y., Lau, R.W.: Content-aware Generative Modeling of Graphic Design Layouts. ACM Transactions on Graphics (TOG) 38(4), 1–15 (2019). https://doi.org/10.1145/3306346.332297
- Zhong, X., Tang, J., Yepes, A.J.: PubLayNet: Largest Dataset Ever for Document Layout Analysis. In: 2019 International Conference on Document Analysis and Recognition (ICDAR). pp. 1015–1022 (2019). https://doi.org/10.1109/ICDAR. 2019.00166
- Zhou, M., Xu, C., Ma, Y., Ge, T., Jiang, Y., Xu, W.: Composition-aware Graphic Layout GAN for Visual-Textual Presentation Designs. In: Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence. pp. 4995– 5001 (2022). https://doi.org/10.24963/ijcai.2022/692