

# LayoutDETR: Detection Transformer Is a Good Multimodal Layout Designer

Ning Yu, Chia-Chih Chen, Zeyuan Chen, Rui Meng, Gang Wu, Paul Josel,  
Juan Carlos Niebles, Caiming Xiong, and Ran Xu

Salesforce Research

{ning.yu, chiachih.chen, zeyuan.chen, ruimeng, gang.wu, pjosel,  
jniebles, cxiong, ran.xu}@salesforce.com

**Abstract.** Graphic layout designs play an essential role in visual communication. Yet handcrafting layout designs is skill-demanding, time-consuming, and non-scalable to batch production. Generative models emerge to make design automation scalable but it remains non-trivial to produce designs that comply with designers’ multimodal desires, i.e., constrained by background images and driven by foreground content. We propose *LayoutDETR* that inherits the high quality and realism from generative modeling, while reformulating content-aware requirements as a detection problem: we learn to detect in a background image the reasonable locations, scales, and spatial relations for multimodal foreground elements in a layout. Our solution sets a new state-of-the-art performance for layout generation on public benchmarks and on our newly-curated ad banner dataset. We integrate our solution into a graphical system that facilitates user studies, and show that users prefer our designs over baselines by significant margins. Code, models, dataset, and demos are available at GitHub.

## 1 Introduction

Graphic layout designs are at the foundation of communication between media designers and their target audience [41, 54, 69]. Multimodal elements, i.e., foreground images/texts, are framed by layout bounding boxes and reasonably arranged on a background image. This relies on a thoughtful understanding of the semantics of each element and their harmony as a whole. Therefore, handcrafting such layout designs is skill-demanding, time-consuming, and requires experienced professionals. In practice, it has been impossible to batch-produce them in massive quantities [9].

Growing demands of automating graphic layout designs have motivated researchers to adapt deep generative models [10, 16, 37, 42, 63] to this task [5, 11, 17, 18, 25, 28, 36, 59, 82, 84], but few of them investigate generation conditioned on multimodal inputs and their fusion for layout designs.

Conditioning on multimodal inputs is critical to enrich designers’ control and to command the aesthetics of layout designs. In this paper, we focus on them. We propose to equip designers with an AI-empowered system that allows multimodal input: arbitrary background images, foreground images, and foreground



**Fig. 1:** **Left:** LayoutDETR takes a background image and a set of multimodal foreground elements (images/texts) as input, and outputs an aesthetically appealing layout. **Right:** we show a few banner samples with rendered texts using our auto-designed layouts. “C” is the composition and “R” the rendering process.

copywriting texts from varying categories. Fig. 1 depicts the functionality and resulting design samples of our solution. We learn a generative model of conditional layout distribution that is constrained by background images and driven by foreground elements. This requires our model to (1) learn the prior distribution from large-scale realistic layout samples, (2) understand the appearance and semantics of background images, (3) understand the appearance and semantics of foreground elements, and (4) fuse background and foreground information to generate the layout bounding box parameters of each foreground element.

To tackle (1), we inherit and explore the high realism from three types of generative learning paradigms: generative adversarial networks (GANs) [16, 35], variational autoencoders (VAEs) [37], and VAE-GAN [42]. To handle (2), we reformulate the background conditioned layout generation as a detection problem, considering both problems require visual understanding and optimize for bounding box parameters. We integrate these two seemingly irrelevant techniques and train a DETR-flavored detector [6] as a layout generator. Specifically, we employ the visual transformer encoder and bounding box transformer decoder architectures of DETR and jointly optimize its supervision loss with generative adversarial loss. We hence name our solution *LayoutDETR*. To handle (3), we incorporate a CNN-based image encoder [8] and a BERT-based text encoder [13] for the multimodal foreground inputs, and feed the features as input tokens to DETR transformer decoder. (4) is naturally handled by DETR thanks to its transformer nature: foreground elements interact with each other through the self-attention in the decoder, while background features interact through the cross-attention between the image encoder and layout decoder.

We summarize our **contributions** as: **(1) Method.** We bridge two seemingly irrelevant research domains, layout generation and visual detection, into one framework that solves multimodal graphic layout design constrained by background images and driven by foreground image/text elements. No existing methods can handle all these modalities at once. **(2) Dataset.** We establish a large-scale ad banner dataset with rich semantic annotations including text bounding boxes, text categories, and text content. We benchmark this dataset for graphic

layout generation over six recent methods, three of our solution variants, and conduct ablation study. We will release the dataset. **(3) State-of-the-art performance.** Our solution reaches a new state-of-the-art performance for graphic layout generation in a comprehensive set of six evaluation metrics, which measure the realism, accuracy, and regularity of generated layouts. **(4) Graphical system and user study.** We integrate our solution into a graphical system that scales up layout generation and facilitates user studies. Users prefer our designs over all baselines by significant margins.

## 2 Related Work

**Deep generative models for layout design.** Automating the layout design with high quality and realism gains increasing attention and achieves substantial progress, thanks to the revolutionary advancements of deep generative models, especially generative adversarial networks (GANs) [4, 16, 35], variational autoencoders (VAEs) [37, 42], autoregressive models (ARMs) [10, 58, 64], and diffusion models [23, 63, 68]. Researchers adopt generative models to learn to generate bounding box parameters, in the form of center location, height, width, and optionally depth, for each foreground element in a layout in either graphics or natural scene domains. Orthogonal to learning paradigms, existing methods also investigate a variety of generator architectures including multi-layer perceptron (MLP) [40], convolutional neural networks (CNN) [66], recursive neural networks (RvNN) [67], long short-term memory (LSTM) [19], graph convolutional networks (GCN) [38], transformer [72], etc. Some layout design methods [5, 17, 18, 36, 51, 59, 81, 82, 84] focus on graphics domain only, while others [14, 18, 32] generalize to natural scenes or even 3D data. Some methods [18, 36, 45, 47, 59, 81] allow only bounding box category condition, while others [5, 17, 43, 48, 82, 84] allow richer conditions in multi-modalities in the form of background images, foreground images, or foreground texts, but not all. A comprehensive taxonomy of layout generation methods is in Table 1.

None of the existing multimodal layout generation methods is designed to handle all the background and foreground modalities at once: LayoutGAN+ [48], NDN [43], LayoutDM [28], PLay [11], DLT [45], LayoutDiffusion [81], LDGM [26], LayoutNUWA [71], Planning and Rendering [50] are unaware of background images in their designs. CGL-GAN [84], ICVT [5], DS-GAN [25], RALF [24] are unaware of foreground texts and images in their formulation. ContentGAN [82] does not take layout spatial information for training and does not even use complete background images as input during training. This does not benefit the layout representation, layout regularity, or final quality (e.g., text readability) after rendering foreground elements onto the background. Vinci [17] relies on a finite set of predefined layout candidates to choose background images from a pool of food and beverage domains. Their method is unable to design layouts conditioned on arbitrary backgrounds in open domains, natural or handcrafted, plain or cluttered, like ours. Considering multimodal layout design persists as a challenging problem, we focus our scope on graphic layouts: mobile application UIs and our newly-organized large-scale ad banners.

**Table 1:** A taxonomy of existing and our layout generation methods. We tag for each method its working data domain(s), data modality(ies) of input condition(s), backend generative model type, as well as implementation architecture. Our method and its variants are the only ones that enable full control of varying multimodal conditions and integrate object detection techniques into layout generation.

Method	Domain(s)	Conditioning modality	Generative model	Architecture
LayoutGAN [47]	Graphics	Bbox category	GAN	MLP, CNN
LayoutGAN+ [48]	Graphics	Bbox category, attribute vector	GAN	MLP, CNN
LayoutGAN++ [36]	Graphics	Bbox category	GAN	Transformer
DS-GAN [25]	Graphics	Bbox category, bg images	GAN	CNN, LSTM
House-GAN [55]	Floor plans	Bubble diagram	GAN	CNN, MPN
House-GAN++ [56]	Floor plans	Bubble diagram	GAN	CNN, MPN
LayoutVAE [32]	Natural scenes	Bbox category	VAE	MLP, LSTM
READ [59]	Graphics	Bbox category	VAE	RvNN
Vinci [17]	Graphics	Bbox category, bg/fg images, text	VAE	LSTM
NDN [43]	Graphics	Bbox category, spatial relation	VAE	GCN
VTN [2]	Graphics, natural scenes	Bbox category	VAE	Transformer
CanvasVAE [76]	Graphics	Bbox category, attribute vector	VAE	Transformer
C2F-VAE [31]	Graphics	Bbox category	VAE	Transformer
DeepSVG [7]	Vector graphics	Vector paths, attribute vector, command	VAE	Transformer
ContentGAN [82]	Graphics	Bbox category, fg image, text	VAE + GAN	Transformer
TextLogoLayout [74]	Graphics	Glyph	VAE + GAN	RNN, CNN
LayoutMCL [57]	Graphics	Bbox category	ARM	RNN, CNN
CanvasEmb [75]	Graphics	Bbox category, attribute vector	ARM	Transformer
LayoutTransformer [18]	Graphics, natural scenes, 3D	Bbox category	ARM	Transformer
BLT [39]	Graphics, natural scenes, 3D	Bbox category	ARM	Transformer
UniLayout [29]	Graphics	Bbox category	ARM	Transformer
LayoutFormer++ [30]	Graphics	Bbox category, spatial relation	ARM	Transformer
Parse-Then-Place [51]	Graphics	Text	ARM	LLM, Transformer
LayoutNUWA [71]	Graphics	Bbox category, HTML code	ARM	LLM, Transformer
LayoutGPT [14]	Natural scenes, 3D	Text	ARM	LLM, Transformer
RALF [24]	Graphics	Bbox category, bg image, spatial relation	ARM	CNN, Transformer
LayoutDM [28]	Graphics	Bbox category, spatial relation	Diffusion	Transformer
PLay [11]	Graphics	Bbox category, guideline	Diffusion	Transformer
DLT [45]	Graphics	Bbox category	Diffusion	Transformer
LayoutDiffusion [81]	Graphics	Bbox category	Diffusion	Transformer
LDGM [26]	Graphics	Bbox category, spatial relation	Diffusion	Transformer
Planning and Rendering [50]	Graphics	Fg image, text	Diffusion	Transformer
CGL-GAN [84]	Graphics	Bg image	GAN + DETR	Transformer
ICVT [5]	Graphics	Bg image	VAE + DETR	Transformer
LayoutDETR-GAN (ours)	Graphics	Bbox category, bg/fg images, text	GAN + DETR	Transformer
LayoutDETR-VAE (ours)	Graphics	Bbox category, bg/fg images, text	VAE + DETR	Transformer
LayoutDETR-VAE-GAN (ours)	Graphics	Bbox category, bg/fg images, text	VAE + GAN + DETR	Transformer

**Object detection.** The goal of object detection is to predict a set of bounding boxes and their categories for each object of interest in a query image. Modern detectors address this in one of three ways: Two-stage detectors [15, 21] predict boxes based on proposals; single-stage detectors predict boxes based on anchors [52] or a grid of possible object centers [60]; direct detectors [6, 12] avoid those hand-crafted initial guesses and directly predict absolute parameters of boxes in the query images. Direct detectors [6, 12] inspire us to leverage their powerful image understanding capability, and combine it with modern layout generators, in order to address the problem of multimodal layout generation. Similar in spirit to CGL-GAN [84] or ICVT [5], we take advantage of the transformer encoder-decoder architecture and generalized intersection over union (gIoU) loss [61] in DETR to learn layout distributions from background image contexts. Unlike prior art, our layout outputs are additionally conditioned on foreground image/text elements.

### 3 LayoutDETR

**Problem statement.** A graphic layout sample  $\mathcal{L}$  is represented by a set of  $N$  2D bounding boxes  $\{\mathbf{b}^i\}_{i=1}^N$ . Each  $\mathbf{b}^i$  is parameterized by a vector with

four elements: its center location in a background image  $(y^i, x^i)$ , height  $h^i$ , and width  $w^i$ . In order to handle background images  $\mathbf{B}$  with arbitrary sizes  $(H, W)$ , we normalize box parameters by their image size correspondingly, i.e.,  $\mathcal{L} = \{(y^i/H, x^i/W, h^i/H, w^i/W)\}_{i=1}^N \doteq \{\hat{\mathbf{b}}^i\}_{i=1}^N$ . The multimodal inputs are the background image  $\mathbf{B}$  and a set of  $N$  foreground elements, in the form of either text elements  $\mathcal{T} = \{\mathbf{t}^i\}_{i=1}^M = \{(\mathbf{s}^i, c^i, l^i)\}_{i=1}^M$  or image patches  $\mathcal{P} = \{\mathbf{p}^i\}_{i=1}^K$ , where  $M \geq 0$ ,  $K \geq 0$ ,  $M + K = N$ .  $\mathbf{s}^i$  is a text string,  $c^i$  is the text class from the set  $\{\text{header text}, \text{body text}, \text{disclaimer text}, \text{button text}\}$ , and  $l^i$  is the length of the text string. Each foreground element corresponds to a bounding box in the layout, indicating its location and size. In case there are foreground elements that we are not interested in, e.g., button underlays or embellishments, we leave them as part of the background. Our goal is to learn a layout generator  $G$  that takes a latent noise vector  $\mathbf{z}$  and the multimodal conditions as input, and outputs a realistic and reasonable layout complying with the multimodal control:  $G(\mathbf{z}, \mathbf{B}, \mathcal{T} \cup \mathcal{P}) \mapsto \mathcal{L}_{\text{fake}}$ .

### 3.1 Generative Learning Frameworks

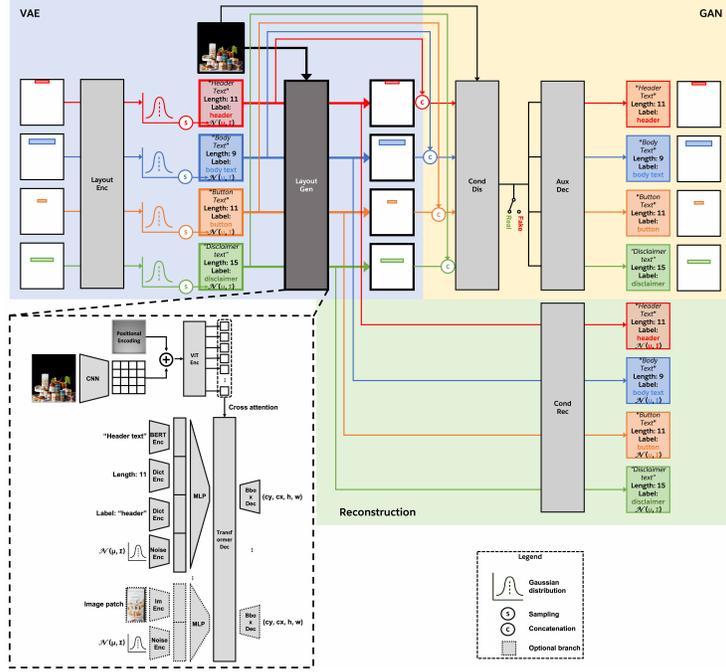
**GAN variant.** Following the GAN paradigm [4, 16, 33–35, 44, 65, 78–80], the generator  $G$  is simultaneously and adversarially trained against discriminator  $D$  training. We formulate a multimodal-conditional discriminator  $D^c(\mathcal{L}, \mathbf{B}, \mathcal{T} \cup \mathcal{P}) \mapsto \{0, 1\}$ , as well as an unconditional discriminator  $D^u(\mathcal{L}) \mapsto \{0, 1\}$ .

It has been observed that the discriminators are insensitive to irregular bounding box positions [36]. For example, the discriminators tend to overlook the unusual layout where a *header texts* is placed at the bottom. As a result, we follow the self-learning technique in [53] and add position-aware regularization to our discriminators. Similar to [36], we add an auxiliary decoder to each discriminator to reconstruct its input. The decoders  $F^c/F^u$  take the output features  $\mathbf{f}^c/\mathbf{f}^u$  of their discriminators  $D^c/D^u$ , add them with learnable positional embeddings  $\mathcal{E}^c = \{\mathbf{e}_i^c\}_{i=1}^N / \mathcal{E}^u = \{\mathbf{e}_i^u\}_{i=1}^N$ , and reconstruct the bounding box parameters and multimodal conditions that are the input of the discriminators:  $F^c(\mathbf{f}^c, \mathcal{E}^c) \mapsto \mathcal{L}_{\text{dec}}^c, \mathbf{B}_{\text{dec}}, \mathcal{T}_{\text{dec}} \cup \mathcal{P}_{\text{dec}} / F^u(\mathbf{f}^u, \mathcal{E}^u) \mapsto \mathcal{L}_{\text{dec}}^u$ . Using  $\mathcal{E}^c/\mathcal{E}^u$  is necessary, without which the reconstructed bounding boxes in a layout would have no variance as they would be conditioned on identical features. The decoders are jointly trained with the discriminators to minimize the reconstruction loss. It enforces the discriminators to fully condition on all their inputs for the binary classification. See Fig. 2 Yellow for the diagram. Thus, our adversarial learning objective is:

$$\min_G \max_{D^c, F^c, \mathcal{E}^c, D^u, F^u, \mathcal{E}^u} L_{\text{GAN}} = L_{\text{GAN\_fake}} + L_{\text{GAN\_real}} \quad (1)$$

$$L_{\text{GAN\_fake}} \doteq \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \{\mathbf{B}, \mathcal{T} \cup \mathcal{P}\} \sim P_{\text{data}}} - \log D^c(\mathcal{L}_{\text{fake}}, \mathbf{B}, \mathcal{T} \cup \mathcal{P}) - \log D^u(\mathcal{L}_{\text{fake}}) \quad (2)$$

$$L_{\text{GAN\_real}} \doteq \mathbb{E}_{\{\mathcal{L}_{\text{real}}, \mathbf{B}, \mathcal{T} \cup \mathcal{P}\} \sim P_{\text{data}}} \log D^c(\mathcal{L}_{\text{real}}, \mathbf{B}, \mathcal{T} \cup \mathcal{P}) + \log D^u(\mathcal{L}_{\text{real}}) - L_{\text{dec}} \quad (3)$$



**Fig. 2:** Our unified training framework covers three generator variants: GAN-, VAE-, and VAE-GAN-based. The layout generator network (darker color and bold) appears in all variants. Its DETR-based multimodal architecture is at the bottom left. During inference, only the generator is needed.

$$L_{\text{dec}} = \lambda_{\text{layout}} (L_{\text{layout}}(\mathcal{L}_{\text{dec}}^{\text{c}}, \mathcal{L}_{\text{real}}) + L_{\text{layout}}(\mathcal{L}_{\text{dec}}^{\text{u}}, \mathcal{L}_{\text{real}})) + \lambda_{\text{im}} (\|\mathbf{B}_{\text{dec}} - \mathbf{B}\|_2 + L_{\text{im}}(\mathcal{P}_{\text{dec}}, \mathcal{P})) + L_{\text{text}}(\mathcal{T}_{\text{dec}}, \mathcal{T}) \quad (4)$$

$$L_{\text{layout}}(\mathcal{L}_1, \mathcal{L}_2) = \frac{1}{N} \sum_{i=1}^N \|\hat{\mathbf{b}}_1^i - \hat{\mathbf{b}}_2^i\|_2 \quad (5)$$

$$L_{\text{im}}(\mathcal{P}_1, \mathcal{P}_2) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{p}_1^i - \mathbf{p}_2^i\|_2 \quad (6)$$

$$L_{\text{text}}(\mathcal{T}_1, \mathcal{T}_2) = \frac{1}{N} \sum_{i=1}^N (\lambda_{\text{str}} L_{\text{str}}(\mathbf{s}_1, \mathbf{s}_2) + \lambda_{\text{cls}} L_{\text{cls}}(c_1, c_2) + \lambda_{\text{len}} L_{\text{len}}(l_1, l_2)) \quad (7)$$

where  $L_{\text{str}}$ ,  $L_{\text{cls}}$ ,  $L_{\text{len}}$  are the reconstruction losses for foreground text strings, text classes, and text lengths respectively.  $L_{\text{str}}$  is the auto-regressive loss according to BERT language modeling [13, 49].  $L_{\text{cls}}$  is the standard classification cross-entropy loss. So is  $L_{\text{len}}$ , considering we quantize and classify a string length integer into one of 256 levels [0, 255].  $\lambda_{\text{layout}} = 500.0$ ,  $\lambda_{\text{im}} = 0.5$ ,  $\lambda_{\text{str}} = 0.1$ ,  $\lambda_{\text{cls}} = 50.0$ , and  $\lambda_{\text{len}} = 2.0$ .

**VAE variant.** VAEs are an alternative paradigm to GANs for generative models. Following the VAE paradigm [37, 62], the generator  $G$  is jointly trained with an encoder  $E$  that maps from the layout space to the latent noise distribution space. The output of  $E$  are the mean  $\mu$  and covariance matrix  $\Sigma$  of a multivariate Gaussian distribution,  $E(\mathcal{L}) \mapsto \mu, \Sigma$ , the samples of which are input to  $G$ :  $\text{sample}(\mu, \Sigma) = \mathbf{z}_0^T \Sigma^{\frac{1}{2}} \mathbf{z}_0 + \mu$ ,  $\mathbf{z}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . It represents the differentiable reparameterization trick in the standard VAE pipeline [37, 62]. See Fig. 2 Blue.

The conditional VAE minimizes the reconstruction loss:

$$\min_{E, G} L_{\text{VAE}} \doteq \mathbb{E}_{\{\mathcal{L}_{\text{real}}, \mathbf{B}, \mathcal{T} \cup \mathcal{P}\} \sim P_{\text{data}}} \lambda_{\text{layout}} L_{\text{layout}}(\mathcal{L}_{\text{fake}}, \mathcal{L}_{\text{real}}) + \lambda_{\text{KL}} \text{KL}(E(\mathcal{L}_{\text{real}}) || \mathcal{N}(\mathbf{0}, \mathbf{I})) \quad (8)$$

$\text{KL}(\cdot || \cdot)$  is the Kullback-Leibler (KL) Divergence used in standard VAE [37, 62] to regularize the encoded latent noise distribution. The hyper-parameters  $\lambda_{\text{layout}} = 500.0$  and  $\lambda_{\text{KL}} = 1.0$ .

**VAE-GAN variant.** VAEs and GANs are also compatible with each other. See Fig. 2 Blue plus Yellow. Following [42, 82], we jointly optimize Eq. 1 and 8:

$$\min_{E, G} \max_{D^c, F^c, \mathcal{E}^c, D^u, F^u, \mathcal{E}^u} L_{\text{GAN}} + L_{\text{VAE}} \quad (9)$$

### 3.2 Additional Objectives

Other losses and regularization terms also play important roles in the generated layout quality. First, we add bounding box supervision as in DETR [6]. We use the generalized intersection over union loss  $\text{gIoU}(\cdot, \cdot)$  [61] between generated layout and its ground truth  $L_{\text{gIoU}}(\mathcal{L}_{\text{fake}}, \mathcal{L}_{\text{real}})$ , where:

$$L_{\text{gIoU}}(\mathcal{L}_1, \mathcal{L}_2) \doteq \lambda_{\text{gIoU}} \frac{1}{N} \sum_{i=1}^N \text{gIoU}(\hat{\mathbf{b}}_1^i, \hat{\mathbf{b}}_2^i) \quad (10)$$

where the hyper-parameter  $\lambda_{\text{gIoU}} = 4.0$ .

Second, we introduce an auxiliary reconstructor  $R$  for the generator to enhance the controllability of input conditions.  $R$  takes the last features of  $G$ ,  $\mathcal{F} = \{\mathbf{f}_i^g\}_{i=1}^N$ , as input tokens before outputting box parameters, and learns to reconstruct  $\mathcal{P}$  and  $\mathcal{T}$ :  $R(\mathcal{F}) \mapsto \mathcal{P}_{\text{rec}}, \mathcal{T}_{\text{rec}}$ .

Jointly with  $G$  training, we learn to minimize:

$$L_{\text{rec}} \doteq \lambda_{\text{im}} L_{\text{im}}(\mathcal{P}_{\text{rec}}, \mathcal{P}) + L_{\text{text}}(\mathcal{T}_{\text{rec}}, \mathcal{T}) \quad (11)$$

with  $L_{\text{im}}$  from Eq. 6 and  $L_{\text{text}}$  from 7. See Fig. 2 Green.

Third, reasonable layout designs typically avoid overlapping between foreground elements. We leverage the overlap loss  $L_{\text{overlap}} \doteq \lambda_{\text{overlap}} L_{\text{overlap}}(\mathcal{L}_{\text{fake}})$  from [48] that discourages overlapping between any pair of bounding boxes in a generated layout. We set  $\lambda_{\text{overlap}} = 7.0$ .

Fourth, aesthetically appealing layouts usually maintain one of the six alignments between a pair of adjacent bounding boxes: left, horizontal-center, right, top, vertical-center, and bottom aligned. We leverage the misalignment loss  $L_{\text{misalign}} \doteq \lambda_{\text{misalign}} L_{\text{misalign}}(\mathcal{L}_{\text{fake}})$  follow [48] that discourages misalignment. We set  $\lambda_{\text{misalign}} = 17.0$ .

Finally, our training objective is formulated as follows.

$$\min_{E,G,R} \max_{D^c, F^c, \mathcal{E}^c, D^u, F^u, \mathcal{E}^u} L_{\text{GAN}} + L_{\text{VAE}} + L_{\text{gIoU}} + L_{\text{rec}} + L_{\text{overlap}} + L_{\text{misalign}} \quad (12)$$

All the  $\lambda$ s are trivially set to align the order of magnitude of each loss term. We use an identical set of  $\lambda$ s for all the datasets to validate our performance is insensitive to  $\lambda$ s.

### 3.3 DETR-based Multimodal Architectures

Architecture design is where we integrate object detection with layout generation. Detection transformer (DETR) [6] is employed and modified for Layout-DETR generator  $G$  and conditional discriminator  $D^c$ .

As depicted in Fig. 2 bottom left,  $G$  and  $D^c$  contain a background encoder and a layout transformer decoder. The encoder is the same as in DETR [6]. The decoder is inherited from the DETR decoder with self-attention and encoder-decoder-cross-attention mechanisms [72]. Different from DETR, we replace their freely-learnable object queries with our foreground embeddings as the input tokens to the decoder. The decoder then transforms the embeddings to layout bounding box parameters. Foreground embedding is a concatenation of noise embedding and text/image embedding. Text embedding is a concatenation of text string embedding (BERT-pretrained and fixed), text class embedding (via learning dictionary), and text length embedding (via learning dictionary). See more clarifications on our design choices in the supplementary material. See the implementations of the other networks in the supplementary material. In particular, other networks are transformer-based as well. Only the input and output formats differ depending on the network functionality. We do not choose a wire-frame discriminator since its empirical performance is worse, as also observed in [36].

### 3.4 Motivations of Using Each Network Component

The generator  $G$  and conditional/unconditional discriminators  $D^c/D^u$  serve as the fundamental components for the GAN variant of our solution. The use of conditional discriminator  $D^c$  is straightforward due to the nature of our multimodal conditions: it encourages foreground elements to be harmonic and reasonable to the background after being overlaid according to the generated layout. The unconditional discriminator  $D^u$  differs from  $D^c$  by ignoring background and foreground elements, and focusing only on the realism of layouts themselves. Therefore, the use of  $D^u$  additionally encourages the mutual relations among bounding boxes in a generated layout to be realistic and reasonable, regardless of multimodal conditions. The empirical effectiveness of  $D^u$  is evidenced in Table 2 Row 3.

The motivation of using auxiliary decoders  $F^c/F^u$  following  $D^c/D^u$  is inspired by [36].  $F^c/F^u$  targets to self-reconstruct all the input information of  $D^c/D^u$  through the bottleneck discriminator features. It encourages the input

information to be fully encoded into the discriminator features such that the discriminator classification is fully justified by the input. The empirical effectiveness of  $F^c/F^u$  has been validated in [36] Table 2 last column.

The layout encoder  $E$  and generator  $G$  serve as the fundamental components for the VAE variant of our solution. VAEs and GANs are alternative paradigms of generative models, and are complementary to each other in terms of data distribution learning and feature representation learning [42]. Inspired by [82], we combine the strengths of VAEs and GANs and apply them for layout generation.

The use of auxiliary reconstructor  $R$  following  $G$  stems from the same motivation as the use of auxiliary decoders  $F^c/F^u$ . Its empirical effectiveness is evidenced in Table 2 Row 2.

The motivation of using DETR architecture [6] for background image encoding and understanding stems from its state-of-the-art performance for object detection using the state-of-the-art visual transformer (ViT) encoder architecture [20, 73]. In our scenarios, “detection” is equivalent to “generation” as both processes output bounding box parameters. Although “detection” lacks the “layout” concept, it is complemented by the layout discriminator and encoder networks. “Object” stands for the non-clutter subregions in a background that are suitable for overlaying foreground elements. ViT tokenizes visual features that facilitate cross attention with other foreground element features, e.g., tokenized text features, so that multimodal conditions can synergize jointly.

Incorporating DETR into multimodal layout generation is non-trivial. Directly applying the DETR architecture in the encoder did not achieve optimal performance. Tuning around different generator paradigms (Sec. 3.1), loss configurations and conditional embedding configurations (Sec. 5.1) leads us towards the empirical optimum.

## 4 New Ad Banner Dataset

Not all existing datasets are suitable for multimodal layout design because they do not always provide multi-modality information, e.g. natural scene datasets or Crello graphic document dataset [76]. Other datasets, such as PubLayNet document dataset [83] and PartNet 3D shape dataset [77] render layouts only on plain background. Magazine dataset [82] does not provide complete background images since they have masked out foreground layouts from the background. On the other hand, ad banner datasets are composed of multimodal elements and lead to several previous layout design techniques [5, 17, 25, 84]. Unfortunately, none of their datasets is publicly available, except for CGL dataset [84], PKU PosterLayout [25], TextLogoLayout dataset [74] which, however, do not contain the widely used English modality. We therefore collect a new large-scale ad banner dataset of 7,196 samples containing English characters, which are ready to release.

Each of our samples consists of a well-designed banner image, its layout ground truth, foreground text strings, text classes, and background image. The banner images are filtered from Pitt Image Ads Dataset [27] and crawled from Google Image Search Engine. Their layouts and text classes are manually annotated by Amazon Mechanical Turk (AMT). The text classes are  $\{header\ text,$

*body text, disclaimer text, button text*}, with logos categorized as *header text*. The text strings are extracted by OCR [1] and removed by image inpainting [70] to obtain the text-free background image. Examples and more details about data collection are in the supplementary material.

## 5 Experiments

**Datasets.** Ablation study and user study are conducted on our ad banner dataset with 7,196 samples. Comparisons to baselines are additionally performed on the CGL dataset with 59,978 valid Chinese ad banner samples [84], and on the CLAY dataset with 32,063 valid mobile application UI samples [46]. Besides background image and foreground text conditions, it also involves foreground image conditions. We apply the same OCR and image inpainting processes to CGL and CLAY datasets, in order to extract texts as part of input conditions, and separate apart foreground from background. 90% of the samples are used for training and 10% for testing.

**Baselines.** We select six recent methods covering a variety of generative paradigms and architectures listed in Table 1: LayoutGAN++ [36], READ [59], Vinci [17], LayoutTransformer [18], CGL-GAN [84], and ICVT [5]. We noted that LayoutTransformer [18] is empirically superior to the more recent work BLT [39], the same observation as Row 1 v.s. Row 8 in Table 2 of [29]. We therefore did not experiment with BLT. Several baselines do not allow background conditions and/or foreground text/image conditions. We integrate our encoders to their models for fair comparisons.

**Evaluation metrics.** (1) For the **realism** of generated layouts, we calculate the Fréchet distances [22] and kernel distances [3] between fake and real feature distributions. All the real testing samples are used, and the same number of generated samples are used. We consider two feature spaces: the layout features pretrained by [36], and VGG image features pretrained on ImageNet [22]. We obtain the output banner images by overlaying foreground image patches and rendering foreground texts on top of background images according to the generated layout. The rendering process and examples are in Fig. 4. (2) For the sample-wise **accuracy** of generated layouts w.r.t. their ground truth, we calculate the layout-to-layout IoU [36] and DocSim [59]. Box-level matching is trivial as the correspondences between generated and ground truth boxes are given by the conditioning input. (3) For the **regularity** of generated layouts, our metrics are the overlap loss [48] and misalignment loss [48] in Section 3.2.

### 5.1 Ablation Study

**Loss configurations.** We start from the LayoutGAN++ [36] baseline implementation and additionally enable it to take multimodal foreground and background as input conditions. We then progressively add on extra loss terms and report the quantitative measurements in Table 2 top section. We observe: (1) Row 1 contains the far worse results of randomly generated layouts on real backgrounds, indicating that the quality of layouts itself matters. (2) Comparing Row

**Table 2:** Ablation study w.r.t. loss config (top) or conditional embedding config (bottom) on **our ad banner dataset**. Each row is progressively ablated from its row above. Each cell contains mean $\pm$ std. For FIDs and KIDs, they are the statistics over 10 runs.  $\uparrow/\downarrow$  indicates a higher/lower value is better. **Bold/underline** font indicates the top/second best value in the column.

Method	Realism				Accuracy		Regularity	
	Layout FID $\downarrow$	Layout KID ( $\times 10^{-3}$ ) $\downarrow$	Image FID $\downarrow$	Image KID ( $\times 10^{-5}$ ) $\downarrow$	IoU $\uparrow$	DocSim $\uparrow$	Overlap $\downarrow$	Misalign ( $\times 10^{-2}$ ) $\downarrow$
Random layout on real bg	58.21 $\pm$ 4.04	525.93 $\pm$ 45.08	51.01 $\pm$ 0.41	582.47 $\pm$ 7.53	-	-	-	-
Conditional LayoutGAN++	11.33 $\pm$ 0.10	44.77 $\pm$ 0.36	36.06 $\pm$ 0.02	115.16 $\pm$ 3.37	0.111 $\pm$ 0.001	0.121 $\pm$ 0.001	0.374 $\pm$ 0.006	2.194 $\pm$ 0.058
+ Aux. Dec. (Eq. 4-7)	4.25 $\pm$ 0.01	16.62 $\pm$ 0.05	28.40 $\pm$ 0.06	58.5 $\pm$ 1.45	0.163 $\pm$ 0.002	0.130 $\pm$ 0.001	0.104 $\pm$ 0.003	0.759 $\pm$ 0.021
+ Gen. Rec. (Eq. 11)	3.27 $\pm$ 0.01	11.80 $\pm$ 0.04	29.56 $\pm$ 0.06	11.29 $\pm$ 0.20	0.186 $\pm$ 0.002	<u>0.148</u> $\pm$ 0.001	0.125 $\pm$ 0.003	0.853 $\pm$ 0.016
+ Uncond. Dis. $D^u$	3.70 $\pm$ 0.05	16.23 $\pm$ 0.08	29.21 $\pm$ 0.08	25.09 $\pm$ 0.02	0.177 $\pm$ 0.002	0.140 $\pm$ 0.001	<u>0.103</u> $\pm$ 0.003	<u>0.681</u> $\pm$ 0.011
+ gIoU loss (Eq. 10)	<u>3.23</u> $\pm$ 0.01	11.60 $\pm$ 0.02	<u>28.20</u> $\pm$ 0.04	<u>10.51</u> $\pm$ 0.09	0.182 $\pm$ 0.002	0.138 $\pm$ 0.001	0.106 $\pm$ 0.003	0.721 $\pm$ 0.011
+ Overlap & Misalign loss = LayoutDETR-GAN (ours)	<b>3.19</b> $\pm$ 0.01	<b>5.62</b> $\pm$ 0.01	<b>27.35</b> $\pm$ 0.04	<b>8.31</b> $\pm$ 0.80	<b>0.208</b> $\pm$ 0.002	<b>0.151</b> $\pm$ 0.000	<b>0.101</b> $\pm$ 0.003	<b>0.646</b> $\pm$ 0.011
- Text length embeddings	3.24 $\pm$ 0.01	<u>9.25</u> $\pm$ 0.05	28.65 $\pm$ 0.03	11.42 $\pm$ 0.35	<u>0.191</u> $\pm$ 0.002	0.144 $\pm$ 0.001	0.117 $\pm$ 0.003	0.807 $\pm$ 0.012
- Text class embeddings	25.17 $\pm$ 0.54	171.88 $\pm$ 5.17	29.25 $\pm$ 0.25	139.16 $\pm$ 4.44	0.166 $\pm$ 0.002	0.132 $\pm$ 0.001	0.110 $\pm$ 0.001	0.000 $\pm$ 0.000

2 and 3, all the metrics are improved, as the auxiliary decoder enhances the discriminator’s conditioning and representation. (3) Comparing Row 3 and 4, all the metrics are improved, thanks to the enhanced controllability through the reconstruction of conditional inputs. (4) Comparing Row 4 and 5, unconditional discriminator benefits the layout regularity due to its approximation power between generated layout parameters and real regular ones. (5) Comparing Row 5 and 6, the supervised gIoU loss boosts the realism and accuracy by a significant margin, yet seemingly contradicts the regularity. (6) Fortunately, in Row 7, adding overlap loss and misalignment loss optimizes all the metrics. We name this "LayoutDETR-GAN (ours)" and stick to it for the following experiments. (7) Error margins after " $\pm$ " are consistently smaller than value differences across rows, indicating the differences are statistically significant.

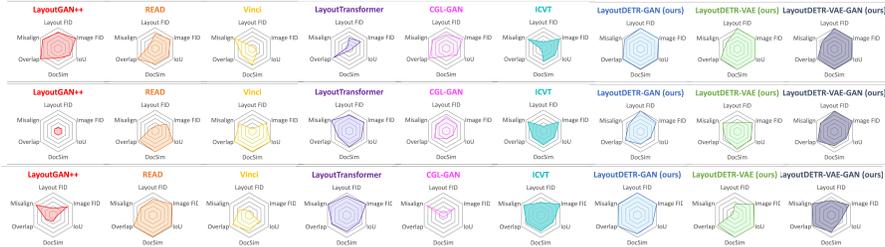
**Conditional embedding configurations.** We do not consider foreground images in this ablation study as the embeddings are simply image features. Whereas for foreground texts, we examine the importance of text length embeddings and text class embeddings by progressively removing them from the training. From Table 2 bottom section we observe: (1) Comparing Row 7 and 8, text length embeddings are beneficial all around. We reason that text length is a strong indicator of a proper text box size. This validates the strong positive correlation between text lengths and box sizes. (2) Comparing Row 8 and 9, text label embeddings serve as an essential role in the generation. Layout FID and Layout KID deteriorate significantly without text label embeddings (Row 9) as bounding boxes of similar texts tend to collapse to the same regions (referring to the 0.0 misalignment). This implies that text content itself is not as discriminative as text labels to differentiate box parameters.

## 5.2 Comparisons to baselines

**Quantitative comparisons** are in Table 3. We evaluate three of our alternative model variants: GAN-, VAE-, and VAE-GAN-based. For intuitive understanding of the tables, we visualize each row into a radar plot in Fig. 3. Each corner of the

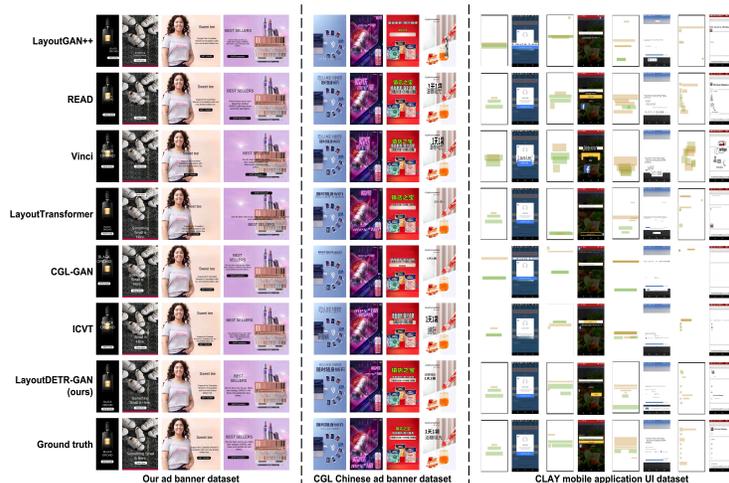
**Table 3:** Quantitative comparisons to baselines. Each cell contains mean $\pm$ std. For FIDs and KIDs, they are the statistics over 10 runs.  $\uparrow/\downarrow$  indicates a higher/lower value is better. **Bold/underline** font indicates the top/second best value in the column.

Method	Realism				Accuracy		Regularity	
	Layout FID $\downarrow$	Layout KID ( $\times 10^{-3}$ ) $\downarrow$	Image FID $\downarrow$	Image KID ( $\times 10^{-5}$ ) $\downarrow$	IoU $\uparrow$	DocSim $\uparrow$	Overlap $\downarrow$	Misalign ( $\times 10^{-2}$ ) $\downarrow$
<b>Our ad banner dataset</b>								
LayoutGAN++	4.25 $\pm$ 0.01	16.62 $\pm$ 0.05	28.40 $\pm$ 0.06	58.54 $\pm$ 1.45	0.163 $\pm$ 0.002	0.130 $\pm$ 0.001	0.104 $\pm$ 0.003	0.759 $\pm$ 0.021
READ	4.45 $\pm$ 0.02	15.21 $\pm$ 0.21	32.10 $\pm$ 0.13	77.53 $\pm$ 2.23	0.177 $\pm$ 0.002	0.141 $\pm$ 0.001	<b>0.093</b> $\pm$ 0.002	2.867 $\pm$ 0.040
Vinci	38.97 $\pm$ 0.10	231.70 $\pm$ 1.22	58.12 $\pm$ 0.20	833.00 $\pm$ 3.55	0.104 $\pm$ 0.001	0.143 $\pm$ 0.001	0.243 $\pm$ 0.003	<b>0.271</b> $\pm$ 0.010
LayoutTransformer	5.47 $\pm$ 0.01	13.87 $\pm$ 0.01	39.70 $\pm$ 0.01	134.87 $\pm$ 1.03	0.080 $\pm$ 0.001	0.115 $\pm$ 0.001	0.127 $\pm$ 0.003	3.632 $\pm$ 0.065
CGL-GAN	4.69 $\pm$ 0.01	17.58 $\pm$ 0.02	30.50 $\pm$ 0.02	13.52 $\pm$ 1.40	0.154 $\pm$ 0.002	0.127 $\pm$ 0.001	0.116 $\pm$ 0.003	1.191 $\pm$ 0.025
ICVT	12.54 $\pm$ 0.06	64.49 $\pm$ 0.12	30.11 $\pm$ 0.05	62.29 $\pm$ 2.54	0.163 $\pm$ 0.002	0.137 $\pm$ 0.001	0.423 $\pm$ 0.006	0.682 $\pm$ 0.018
LayoutDETR-GAN	<b>3.19</b> $\pm$ 0.01	<b>5.62</b> $\pm$ 0.01	<b>27.35</b> $\pm$ 0.04	8.31 $\pm$ 0.80	0.208 $\pm$ 0.002	<u>0.151</u> $\pm$ 0.000	0.101 $\pm$ 0.003	0.646 $\pm$ 0.011
LayoutDETR-VAE	3.25 $\pm$ 0.03	11.97 $\pm$ 0.26	<u>27.47</u> $\pm$ 0.04	<u>7.70</u> $\pm$ 0.22	<b>0.216</b> $\pm$ 0.002	<b>0.152</b> $\pm$ 0.001	0.119 $\pm$ 0.002	1.737 $\pm$ 0.037
LayoutDETR-VAE-GAN	<u>3.23</u> $\pm$ 0.02	<u>10.75</u> $\pm$ 0.09	27.88 $\pm$ 0.11	<b>4.18</b> $\pm$ 0.24	<u>0.210</u> $\pm$ 0.002	<u>0.151</u> $\pm$ 0.001	0.117 $\pm$ 0.002	1.439 $\pm$ 0.023
<b>CGL Chinese ad banner dataset</b>								
LayoutGAN++	11.43 $\pm$ 0.05	59.02 $\pm$ 0.26	11.92 $\pm$ 0.05	1082.68 $\pm$ 20.71	0.061 $\pm$ 0.001	0.083 $\pm$ 0.000	0.593 $\pm$ 0.007	0.729 $\pm$ 0.017
READ	10.51 $\pm$ 0.04	107.30 $\pm$ 1.24	6.58 $\pm$ 0.08	465.96 $\pm$ 14.92	<b>0.269</b> $\pm$ 0.002	<b>0.127</b> $\pm$ 0.001	0.145 $\pm$ 0.002	0.704 $\pm$ 0.098
Vinci	12.06 $\pm$ 0.01	80.45 $\pm$ 0.43	5.38 $\pm$ 0.02	320.99 $\pm$ 6.30	<u>0.266</u> $\pm$ 0.002	<u>0.125</u> $\pm$ 0.001	<b>0.093</b> $\pm$ 0.002	0.433 $\pm$ 0.042
LayoutTransformer	5.11 $\pm$ 0.01	33.72 $\pm$ 0.54	4.65 $\pm$ 0.01	286.08 $\pm$ 2.99	0.186 $\pm$ 0.002	0.114 $\pm$ 0.001	0.340 $\pm$ 0.005	<b>0.276</b> $\pm$ 0.027
CGL-GAN	5.63 $\pm$ 0.01	36.99 $\pm$ 0.30	7.26 $\pm$ 0.09	744.49 $\pm$ 12.73	0.107 $\pm$ 0.001	0.093 $\pm$ 0.001	0.297 $\pm$ 0.004	0.538 $\pm$ 0.011
ICVT	10.76 $\pm$ 0.14	119.10 $\pm$ 1.14	4.22 $\pm$ 0.05	109.33 $\pm$ 3.85	0.169 $\pm$ 0.002	0.109 $\pm$ 0.001	0.327 $\pm$ 0.004	<u>0.340</u> $\pm$ 0.060
LayoutDETR-GAN	<b>2.40</b> $\pm$ 0.01	<u>13.60</u> $\pm$ 0.29	<u>4.11</u> $\pm$ 0.06	<u>22.60</u> $\pm$ 0.79	0.157 $\pm$ 0.002	0.106 $\pm$ 0.000	0.187 $\pm$ 0.003	0.464 $\pm$ 0.010
LayoutDETR-VAE	8.57 $\pm$ 0.10	94.84 $\pm$ 1.09	4.21 $\pm$ 0.03	144.27 $\pm$ 5.18	0.208 $\pm$ 0.002	0.120 $\pm$ 0.001	0.288 $\pm$ 0.003	0.374 $\pm$ 0.064
LayoutDETR-VAE-GAN	<u>2.65</u> $\pm$ 0.05	<b>12.37</b> $\pm$ 0.60	<b>2.66</b> $\pm$ 0.02	<b>19.30</b> $\pm$ 1.04	0.180 $\pm$ 0.002	0.110 $\pm$ 0.001	<u>0.134</u> $\pm$ 0.002	0.401 $\pm$ 0.070
<b>CLAY mobile application UI dataset</b>								
LayoutGAN++	14.12 $\pm$ 0.06	60.20 $\pm$ 0.52	7.49 $\pm$ 0.02	148.33 $\pm$ 4.02	0.049 $\pm$ 0.001	0.078 $\pm$ 0.001	0.817 $\pm$ 0.115	1.057 $\pm$ 0.028
READ	3.68 $\pm$ 0.01	21.98 $\pm$ 0.15	5.38 $\pm$ 0.02	91.43 $\pm$ 3.24	<u>0.312</u> $\pm$ 0.003	<u>0.121</u> $\pm$ 0.001	<u>0.099</u> $\pm$ 0.002	2.045 $\pm$ 0.045
Vinci	22.98 $\pm$ 0.02	216.90 $\pm$ 0.66	13.04 $\pm$ 0.05	677.42 $\pm$ 7.22	0.178 $\pm$ 0.002	0.104 $\pm$ 0.001	0.253 $\pm$ 0.004	2.526 $\pm$ 0.055
LayoutTransformer	<u>2.64</u> $\pm$ 0.01	<u>5.03</u> $\pm$ 0.18	5.27 $\pm$ 0.01	<u>55.99</u> $\pm$ 2.49	0.216 $\pm$ 0.003	0.106 $\pm$ 0.002	0.357 $\pm$ 0.006	0.833 $\pm$ 0.032
CGL-GAN	47.74 $\pm$ 0.02	190.96 $\pm$ 1.11	8.96 $\pm$ 0.02	226.81 $\pm$ 4.65	0.034 $\pm$ 0.001	0.066 $\pm$ 0.000	1.153 $\pm$ 0.141	1.099 $\pm$ 0.011
ICVT	4.56 $\pm$ 0.04	18.35 $\pm$ 0.26	<u>5.26</u> $\pm$ 0.03	69.83 $\pm$ 2.24	0.208 $\pm$ 0.002	0.105 $\pm$ 0.001	0.396 $\pm$ 0.006	1.066 $\pm$ 0.043
LayoutDETR-GAN	<b>1.84</b> $\pm$ 0.02	<b>3.01</b> $\pm$ 0.15	<b>5.22</b> $\pm$ 0.02	<b>11.19</b> $\pm$ 1.39	0.261 $\pm$ 0.003	0.113 $\pm$ 0.001	<b>0.083</b> $\pm$ 0.002	<u>0.773</u> $\pm$ 0.016
LayoutDETR-VAE	4.99 $\pm$ 0.01	30.18 $\pm$ 0.32	5.49 $\pm$ 0.03	107.55 $\pm$ 3.57	<b>0.327</b> $\pm$ 0.003	<b>0.123</b> $\pm$ 0.001	0.205 $\pm$ 0.004	5.119 $\pm$ 0.019
LayoutDETR-VAE-GAN	3.98 $\pm$ 0.12	18.39 $\pm$ 0.98	5.87 $\pm$ 0.02	82.31 $\pm$ 1.56	0.158 $\pm$ 0.002	0.108 $\pm$ 0.001	0.148 $\pm$ 0.002	<b>0.691</b> $\pm$ 0.030



**Fig. 3:** Radar plots for Table 3 on **our ad banner dataset (top)**, **CGL Chinese ad banner dataset (middle)**, and **CLAY mobile application UI dataset (bottom)**. Each plot corresponds to a row (method) in the table. Each corner in a plot corresponds to a column (metric) in the table. Values are normalized to the unit range, the higher the better.

radar plot corresponds to a metric. We observe: (1) On our ad banner dataset, at least one of our variants outperforms all the baselines in terms of realism and accuracy. Our LayoutDETR-GAN achieves the second best in terms of regularity. The margin to the best baseline is minor. This evidences the efficacy of our understanding of multi-modality and sophisticated loss configurations. (2) On CGL Chinese ad banner dataset, our variants lead in the most important realism metrics by significant margins. Our LayoutDETR-VAE-GAN variant is the



**Fig. 4:** **Left:** comparisons on the testing set of **our ad banner dataset**. We apply the same rendering process to all methods: (1) Text font sizes and line breakers are adaptively determined to tightly fit into their inferred boxes. (2) Text font colors and button pad colors are adaptively determined to be either black or white whichever contrasts more with the background. (3) Button text colors are then determined to contrast with the button pads. (4) Text font is set to *Arial*. (5) Boxes are enforced to horizontally center-align with each other. **Middle:** comparisons on **CGL Chinese ad banner dataset**. Image patches that contain foreground text elements are resized and overlaid on the background following the generated layouts. **Right:** comparisons on **CLAY mobile application UI dataset**.

best overall and outperforms the LayoutDETR-GAN variant. This is because CGL has more elements per sample and VAE plays a more important role than GAN for complex layout arrangements. LayoutTransformer achieves a similarly balanced performance. READ and Vinci lead in the accuracy metrics yet significantly underperform in at least one other metric. (3) On CLAY dataset, at least one of our variants outperforms all the baselines in all metrics. Our efficacy generalizes in at least these two multimodal foreground domains: texts and images. (4) Error margins after “ $\pm$ ” are consistently smaller than value differences across rows, indicating the differences are statistically significant.

**Qualitative comparisons** are in Fig. 4. We observe: (1) For our and CGL ad banner datasets, our designs understand the background the most effectively. For example, they never overlay foreground elements on top of clutter background subregions. If the background looks symmetric, our layouts are placed in the middle of the banners. (2) For these two datasets, our designs also approximate the real layout distribution the most closely, in terms of the relative font sizes (box sizes) and spatial relations (box orders and distances). Even for samples not close to their ground truth, our design variants still look the most reasonable and most harmonic together with backgrounds. (3) CLAY dataset is more challenging in terms of multimodal conditioning, as it has more tiny foreground elements. Still, our layouts appear the best aligned and least overlapping, with reasonable designs to harmonize with backgrounds, although different from

**Table 4:** Pairwise user preferences (column method over row method) on our ad banner dataset.

Method	READ	Vinci	LayoutTransformer	CGL-GAN	ICVT	LayoutDETR-GAN (ours)
LayoutGAN++	49.8% <sub>p=0.4</sub>	45.6% <sub>p=3e-3</sub>	44.4% <sub>p=3e-4</sub>	53.9% <sub>p=0.01</sub>	47.1% <sub>p=0.04</sub>	55.7% <sub>p=2e-4</sub>
READ	-	45.1% <sub>p=1e-3</sub>	44.5% <sub>p=3e-4</sub>	53.8% <sub>p=0.01</sub>	53.0% <sub>p=0.04</sub>	54.2% <sub>p=5e-3</sub>
Vinci	-	-	51.7% <sub>p=0.2</sub>	55.8% <sub>p=2e-4</sub>	56.9% <sub>p=1e-5</sub>	62.6% <sub>p=3e-15</sub>
LayoutTransformer	-	-	-	57.1% <sub>p=8e-6</sub>	56.0% <sub>p=2e-4</sub>	63.5% <sub>p=2e-17</sub>
CGL-GAN	-	-	-	-	48.9% <sub>p=0.2</sub>	54.7% <sub>p=3e-3</sub>
ICVT	-	-	-	-	-	55.4% <sub>p=6e-4</sub>

the ground truth. The supplementary material and attached video show more uncurated results, including the impact of varying texts on layouts and our limitation in challenging scenarios.

### 5.3 Graphical System Design and User Study

We integrate generators into a graphical system for user-friendly applications in practice. The UI design and demo video are attached in the supplementary material. With the graphical system, we can test a massive number of cases and collect generation results in the wild. This facilitates us to analyze users’ subjective preferences for our designs. In specific, we tested on 308 ad banner samples by rendering our designs and baseline designs via the graphical system. We configured a binary comparison task between each pair of designs. In total, we launched  $308 \times \binom{7}{2} = 6,468$  jobs on AMT, and randomly assigned three workers for each job. We simply asked each worker “Which of the two images looks better?” We intentionally did not pre-define any criterion for the preference so as to fully respect their subjective judgments. All workers have an approval rate history above 90% on AMT. We did not set any restrictions for workers’ gender, race, sexuality, demographics, locations, remuneration rates, etc. Our user study has been reviewed and approved by our ethical board.

Table 4 lists the ratio of users that prefer one design over another. To quantify the statistical significance of our user study, we calculate the p-value of a null hypothesis that the results of binary comparisons are equivalent to tossing a fair coin. We calculate it via the cumulative distribution function of binomial distribution, the smaller the more significant. In the last column all ratios are above 50%: a majority of users prefer our designs over any baseline significantly, considering all p-values  $\ll 0.05$ . This validates that our layout designs are more appealing to users. We attribute this to our effective approximation of real layout distributions and multimodal understanding.

## 6 Conclusion

We present LayoutDETR for customizable layout design. It inherits the high quality and realism of generative models, and benefits from object detectors to understand multimodal conditions. Experiments show that we achieve a new state-of-the-art performance for layout generation on a new ad banner dataset and beyond. We implement our solution as a graphical system that facilitates user studies. Users prefer our designs over several recent works.

## Acknowledgments

We thank Shu Zhang, Silvio Savarese, Abigail Kuttruff, Brian Brechbuhl, Elham Etemad, and Amrutha Krishnan from Salesforce for constructive advice.

## References

1. <https://github.com/PaddlePaddle/PaddleOCR>
2. Arroyo, D.M., Postels, J., Tombari, F.: Variational transformer networks for layout generation. CVPR (2021)
3. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd gans. ICLR (2018)
4. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. ICLR (2019)
5. Cao, Y., Ma, Y., Zhou, M., Liu, C., Xie, H., Ge, T., Jiang, Y.: Geometry aligned variational transformer for image-conditioned layout generation. Multi-media (2022)
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. ECCV (2020)
7. Carlier, A., Danelljan, M., Alahi, A., Timofte, R.: Deepsvg: A hierarchical generative network for vector graphics animation. NeurIPS (2020)
8. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments (2020)
9. Chen, G., Xie, P., Dong, J., Wang, T.: Understanding programmatic creative: The role of ai. Journal of Advertising (2019)
10. Chen, X., Mishra, N., Rohaninejad, M., Abbeel, P.: Pixelsnail: An improved autoregressive generative model. ICML (2018)
11. Cheng, C.Y., Huang, F., Li, G., Li, Y.: Play: Parametrically conditioned layout generation using latent diffusion. ICML (2023)
12. Dai, Z., Cai, B., Lin, Y., Chen, J.: Up-detr: Unsupervised pre-training for object detection with transformers. CVPR (2021)
13. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. NAACL (2019)
14. Feng, W., Zhu, W., Fu, T.j., Jampani, V., Akula, A., He, X., Basu, S., Wang, X.E., Wang, W.Y.: Layoutgpt: Compositional visual planning and generation with large language models. NeurIPS (2024)
15. Girshick, R.: Fast r-cnn. ICCV (2015)
16. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. NeurIPS (2014)
17. Guo, S., Jin, Z., Sun, F., Li, J., Li, Z., Shi, Y., Cao, N.: Vinci: an intelligent graphic design system for generating advertising posters. CHI (2021)
18. Gupta, K., Lazarow, J., Achille, A., Davis, L.S., Mahadevan, V., Shrivastava, A.: Layouttransformer: Layout generation and completion with self-attention. ICCV (2021)
19. Haykin, S., Network, N.: A comprehensive foundation. Neural networks (2004)
20. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. CVPR (2022)
21. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. ICCV (2017)

22. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS* (2017)
23. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *NeurIPS* (2020)
24. Horita, D., Inoue, N., Kikuchi, K., Yamaguchi, K., Aizawa, K.: Retrieval-augmented layout transformer for content-aware layout generation. *arXiv* (2023)
25. Hsu, H.Y., He, X., Peng, Y., Kong, H., Zhang, Q.: Posterlayout: A new benchmark and approach for content-aware visual-textual presentation layout. *CVPR* (2023)
26. Hui, M., Zhang, Z., Zhang, X., Xie, W., Wang, Y., Lu, Y.: Unifying layout generation with a decoupled diffusion model. *CVPR* (2023)
27. Hussain, Z., Zhang, M., Zhang, X., Ye, K., Thomas, C., Agha, Z., Ong, N., Kovashka, A.: Automatic understanding of image and video advertisements. *CVPR* (2017)
28. Inoue, N., Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: Layoutdm: Discrete diffusion model for controllable layout generation. *CVPR* (2023)
29. Jiang, Z., Deng, H., Wu, Z., Guo, J., Sun, S., Mijovic, V., Yang, Z., Lou, J.G., Zhang, D.: Unilayout: Taming unified sequence-to-sequence transformers for graphic layout generation. *arXiv* (2022)
30. 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. *CVPR* (2023)
31. Jiang, Z., Sun, S., Zhu, J., Lou, J.G., Zhang, D.: Coarse-to-fine generative modeling for graphic layouts. *AAAI* (2022)
32. Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: Layoutvae: Stochastic scene layout generation from a label set. *ICCV* (2019)
33. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. *ICLR* (2018)
34. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. *CVPR* (2019)
35. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. *CVPR* (2020)
36. Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: Constrained graphic layout generation via latent optimization. *ACM MM* (2021)
37. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *ICLR* (2013)
38. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. *ICLR* (2017)
39. Kong, X., Jiang, L., Chang, H., Zhang, H., Hao, Y., Gong, H., Essa, I.: Blt: Bidirectional layout transformer for controllable layout generation. *ECCV* (2022)
40. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *NeurIPS* (2012)
41. Landa, R.: *Graphic design solutions/robin landa*. Wadsworth, Boston (2010)
42. Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. *ICML* (2016)
43. Lee, H.Y., Jiang, L., Essa, I., Le, P.B., Gong, H., Yang, M.H., Yang, W.: Neural design network: Graphic layout generation with constraints. *ECCV* (2020)
44. Lee, K., Chang, H., Jiang, L., Zhang, H., Tu, Z., Liu, C.: Vitgan: Training gans with vision transformers. *ICLR* (2022)
45. Levi, E., Brosh, E., Mykhailych, M., Perez, M.: Dlt: Conditioned layout generation with joint discrete-continuous diffusion layout transformer. *ICCV* (2023)

46. Li, G., Baechler, G., Tragut, M., Li, Y.: Learning to denoise raw mobile ui layouts for improving datasets at scale. CHI (2022)
47. Li, J., Yang, J., Hertzmann, A., Zhang, J., Xu, T.: Layoutgan: Generating graphic layouts with wireframe discriminators. ICLR (2019)
48. Li, J., Yang, J., Zhang, J., Liu, C., Wang, C., Xu, T.: Attribute-conditioned layout gan for automatic graphic design. TVCG (2020)
49. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. ICML (2022)
50. Li, Z., Li, F., Feng, W., Zhu, H., Liu, A., Li, Y., Zhang, Z., Lv, J., Zhu, X., Shen, J., et al.: Planning and rendering: Towards end-to-end product poster generation. arXiv (2023)
51. Lin, J., Guo, J., Sun, S., Xu, W., Liu, T., Lou, J.G., Zhang, D.: A parse-then-place approach for generating graphic layouts from textual descriptions. ICCV (2023)
52. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. ICCV (2017)
53. Liu, B., Zhu, Y., Song, K., Elgammal, A.: Towards faster and stabilized gan training for high-fidelity few-shot image synthesis. ICLR (2020)
54. Lok, S., Feiner, S.: A survey of automated layout techniques for information presentations. Proceedings of SmartGraphics (2001)
55. Nauata, N., Chang, K.H., Cheng, C.Y., Mori, G., Furukawa, Y.: House-gan: Relational generative adversarial networks for graph-constrained house layout generation. ECCV (2020)
56. Nauata, N., Hosseini, S., Chang, K.H., Chu, H., Cheng, C.Y., Furukawa, Y.: House-gan++: Generative adversarial layout refinement network towards intelligent computational agent for professional architects. CVPR (2021)
57. Nguyen, D.D., Nepal, S., Kanhere, S.S.: Diverse multimedia layout generation with multi choice learning. Multimedia (2021)
58. Van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al.: Conditional image generation with pixelcnn decoders. NeurIPS (2016)
59. Patil, A.G., Ben-Eliezer, O., Perel, O., Averbuch-Elor, H.: Read: Recursive autoencoders for document layout generation. CVPR Workshops (2020)
60. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. CVPR (2016)
61. Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. CVPR (2019)
62. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. ICML (2014)
63. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. CVPR (2022)
64. Salimans, T., Karpathy, A., Chen, X., Kingma, D.P.: Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. ICLR (2017)
65. Sauer, A., Schwarz, K., Geiger, A.: Stylegan-xl: Scaling stylegan to large diverse datasets. SIGGRAPH (2022)
66. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. ICLR (2015)
67. Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C.D., Ng, A.Y., Potts, C.: Recursive deep models for semantic compositionality over a sentiment treebank. EMNLP (2013)

68. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. ICLR (2021)
69. Stribley, M.: Rules of composition all designers live by. Retrieved May (2016)
70. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. WACV (2022)
71. Tang, Z., Wu, C., Li, J., Duan, N.: Layoutnuwa: Revealing the hidden layout expertise of large language models. ICLR (2024)
72. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NeurIPS (2017)
73. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. CVPR (2018)
74. Wang, Y., Pu, G., Luo, W., Wang, Y., Xiong, P., Kang, H., Lian, Z.: Aesthetic text logo synthesis via content-aware layout inferring. CVPR (2022)
75. Xie, Y., Huang, D., Wang, J., Lin, C.Y.: Canvasemb: learning layout representation with large-scale pre-training for graphic design. Multimedia (2021)
76. Yamaguchi, K.: Canvasvae: learning to generate vector graphic documents. ICCV (2021)
77. Yu, F., Liu, K., Zhang, Y., Zhu, C., Xu, K.: Partnet: A recursive part decomposition network for fine-grained and hierarchical shape segmentation. CVPR (2019)
78. Yu, J., Li, X., Koh, J.Y., Zhang, H., Pang, R., Qin, J., Ku, A., Xu, Y., Baldrige, J., Wu, Y.: Vector-quantized image modeling with improved vqgan. ICLR (2022)
79. Yu, N., Li, K., Zhou, P., Malik, J., Davis, L., Fritz, M.: Inclusive gan: Improving data and minority coverage in generative models. ECCV (2020)
80. Yu, N., Liu, G., Dundar, A., Tao, A., Catanzaro, B., Davis, L.S., Fritz, M.: Dual contrastive loss and attention for gans. ICCV (2021)
81. Zhang, J., Guo, J., Sun, S., Lou, J.G., Zhang, D.: Layoutdiffusion: Improving graphic layout generation by discrete diffusion probabilistic models. ICCV (2023)
82. Zheng, X., Qiao, X., Cao, Y., Lau, R.W.: Content-aware generative modeling of graphic design layouts. TOG (2019)
83. Zhong, X., Tang, J., Yepes, A.J.: Publaynet: largest dataset ever for document layout analysis. ICDAR (2019)
84. Zhou, M., Xu, C., Ma, Y., Ge, T., Jiang, Y., Xu, W.: Composition-aware graphic layout gan for visual-textual presentation designs. IJCAI (2022)