

Aesthetic Text Logo Synthesis via Content-aware Layout Inferring

Yizhi Wang^{1*}, Guo Pu¹, Wenhan Luo², Yexin Wang², Pengfei Xiong², Hongwen Kang², Zhouhui Lian^{1†}
¹Wangxuan Institute of Computer Technology, Peking University, China
²PCG, Tencent

Abstract

Text logo design heavily relies on the creativity and expertise of professional designers, in which arranging element layouts is one of the most important procedures. However, few attention has been paid to this task which needs to take many factors (e.g., fonts, linguistics, topics, etc.) into consideration. In this paper, we propose a content-aware layout generation network which takes glyph images and their corresponding text as input and synthesizes aesthetic layouts for them automatically. Specifically, we develop a dual-discriminator module, including a sequence discriminator and an image discriminator, to evaluate both the character placing trajectories and rendered shapes of synthesized text logos, respectively. Furthermore, we fuse the information of linguistics from texts and visual semantics from glyphs to guide layout prediction, which both play important roles in professional layout design. To train and evaluate our approach, we construct a dataset named as TextLogo3K, consisting of about 3,500 text logo images and their pixel-level annotations. Experimental studies on this dataset demonstrate the effectiveness of our approach for synthesizing visually-pleasing text logos and verify its superiority against the state of the art.

1. Introduction

Creative AI has attracted considerable interests in recent years. Automatic layout design is a promising and valuable technique, which is quite helpful to facilitate the tasks of designers and improve the efficiency of creating media content. One intuitive way to achieve this goal is to generate layouts by defining specific easy-to-execute rules, but the results are typically plain and without design aesthetics. Therefore, various methods [1, 5, 6, 11, 15–17, 35, 38] based on deep generative models (such as GAN [4] and VAE [12]) have been proposed recently to synthesize graphic layouts by learning from human-designed data. Text logo design is a challenging task because many factors need to be taken



Figure 1. Examples of some commonly seen layout types for text logos. “Trans.” denotes transformation. Yellow boxes denote the bounding boxes of characters/words.

into consideration, such as fonts, layouts, and textures. As we know, the layout of a text logo is closely related to the logo’s text and the selected font styles. However, existing approaches [5, 16, 17, 38] exploit only the visual, category, and topic (keyword) information of elements, or combine some of them to generate the layouts. In fact, the semantics of texts are also very important to determine the layout design. As shown in Figure 1, in a desired layout, the Chinese character “爱” (“love”) is a verb and the key character in the text logo, so it is emphasized by a larger font size than other characters such as “如果” (“if”), which is a less-important conjunction word. Moreover, the synthesized layouts need to take fine-grained details into consideration, such as avoiding the collision of strokes of different glyphs. In addition, the placing trajectories of characters should follow a correct order for reading (e.g., left to right and up to bottom for English) and possess diverse styles at the same time, which cannot be easily handled by non-sequence generation models.

In this paper, we cast the layout synthesis of text logos as a sequence generation problem and employ a GAN-based model¹ to solve it. In our approach, a dual-modality fusion scheme is proposed to encode both the visual and linguistic cues of the input glyphs/text as the condition of GAN. It is widely acknowledged that using GAN to generate discrete sequences is difficult. However, the geometric parameters of design elements are continuous in our task, making GAN training on the sequence of geometric parameters

*This work was partly done when Yizhi Wang was an intern in Tencent

†Corresponding author. E-mail: lianzhouhui@pku.edu.cn

¹Code and dataset: <https://github.com/yizhiwang96/TextLogoLayout>

feasible. As mentioned above, fine-grained details such as glyph collisions and character placing trajectories are critical for the aesthetic quality of a text logo. To this end, a dual-discriminator module is proposed to capture both the character placing trajectory and rendered shape of the synthesized text logo. The placing trajectories of characters are supervised by a sequence discriminator, to guarantee that (1) they follow the correct reading order; and (2) they possess various styles like the human-designed ones. Moreover, by proposing a differentiable composition method, we build a bridge between the synthesized layouts and rendered text logo images. At the same time, an image discriminator is also employed to capture the detailed structures of the synthesized text logos, whose visual quality can be further improved in this manner.

To train and evaluate our model, a large-scale dataset of text logos, TextLogo3K, is constructed with pixel-wise annotations. Extensive experiments including ablation studies and method comparison have been conducted on the new dataset. Experimental results clearly demonstrate the effectiveness and advantages of our proposed method. We also demonstrate how our proposed model can be integrated with font generation [33] and texture transfer [22] models to automatically synthesize visually-pleasing text logos. To the best of our knowledge, this work presents the first text logo synthesis method that is capable of taking both linguistic (i.e., text contents) and visual (i.e., glyph shapes, font styles, and textures) information into account like professional human designers. Furthermore, our TextLogo3K is also the first dataset that specifically targets at the task of text logo synthesis, which will be released to promote the investigation of this topic and other relevant tasks in the community.

2. Related Work

GANs. In recent years, Generative Adversarial Networks (GANs) [4] have achieved great success in the area of image synthesis. Mirza et al. [23] proposed Conditional GAN by adding the data label as a parameter to the input of the generator and discriminator to make the network generate the desired data output. GANs have been employed in various tasks like image SR [14], deblurring [13, 24], dehazing [18, 28], text to image translation [9, 27] and video prediction [20, 32].

Graphic Layout Generation. Li et al. [16] proposed a GAN-based model equipped with wireframe discriminators, where Differentiable Wireframe Rendering is utilized to rasterize the input graphic elements into 2D wireframe images. Zheng et al. [38] designed a conditional GAN model to arrange the layouts of paragraphs and images for posters by encoding the input images, keywords, and attributes. Guo et al. [5] developed a VAE-based model for generating the locations of design elements (e.g., background images, embellishments, and text styles). More re-

cently, Lee et al. [15] proposed a GCN-based model to generate layouts according to both the design components and the specified constraints. Li et al. [17] proposed attribute-conditioned layout GAN to synthesize layouts according to user-specified layout attributes (such as expected area, aspect ratio and reading order).

Unfortunately, none of them can properly handle our task. Instead of using the wireframe images [16], we propose to rasterize glyphs directly to the canvas so that the discriminator can capture more details from the inputs (such as strokes). The reading order in [17] is rule-based and predefined while our method generates content-adaptive reading orders. Through experiments, we also find that sequence generative models are more appropriate than CNNs (such as [38]) to generate the location of each character.

Sequence Modeling. Recurrent Neural Networks (RNNs) such as LSTM [8] and GRU [2] have been very successful in NLP. Recently, Transformers [31] become popular and outperform RNNs when trained with huge amounts of data. As suggested in Vinci [5], sequential ordering is important for designing. For our task, the position of each glyph image in a text logo should follow the corresponding reading order. Therefore, our task can be viewed as a sequence generation problem where a set of coordinates is generated for one design element at each step.

Differentiable Composition. Reddy et al. [26] proposed a differentiable composition method which can find the type, position, orientation, and layering of elements that form a given pattern image. However, their method did not deal with the scaling transformation, which is important in layout design. In this paper, we transform each glyph image into the predicted position on the canvas via a variant of Spatial Transform Networks (STN) [10].

3. Method

3.1. Preliminaries

As shown in Figure 2, a text logo is comprised of a set of N glyph elements: $(g_1, p_1), \dots, (g_N, p_N)$, where $g_i \in \mathbb{R}^{H_g \times W_g}$ is the raw glyph image, W_g and H_g are the width and height of the glyph image, respectively, and p_i denotes the geometric parameters of g_i being presented on the canvas. We assume that the bounding boxes of glyphs are rectangles, i.e., $p_i = (x_i^c, y_i^c, w_i, h_i)$, where (x_i^c, y_i^c) are the coordinates of the centre point, w_i and h_i are the width and height of the rectangle, respectively; $x_i^c, w_i \in (0, W_c)$ and $y_i^c, h_i \in (0, H_c)$, where H_c and W_c are the height and width of the logo canvas, respectively.

For our layout generation network, the input consists of the logo's text and its corresponding glyph images $\mathbf{g} = (g_1, \dots, g_N)$. We represent the input text by the character embeddings provided by [19], which are denoted as $\mathbf{f}^e = (f_1^e, f_2^e, \dots, f_N^e)$. The ground-truth (human-designed)

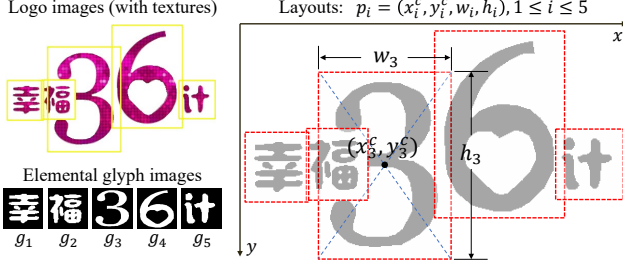


Figure 2. An illustration of the layout and elemental glyph images of a text logo image.

layout is denoted as $\mathbf{p} = (p_1, \dots, p_N)$.

3.2. Overview

We illustrate the pipeline of our layout generation model in Figure 3 and give more details in the following subsections. Specifically, we exploit the dual-modality features of input elements (*i.e.*, visual features \mathbf{f}^v and character embeddings \mathbf{f}^e) and encode them as the condition features \mathbf{f}^c . The **Coordinate Generator** takes the condition vector \mathbf{f}^c and a random noise z as inputs, to predict the geometric parameters of each character that are denoted as $\hat{\mathbf{p}} = (\hat{p}_1, \hat{p}_2, \dots, \hat{p}_N)$, where $\hat{p}_i = (\hat{x}_i^c, \hat{y}_i^c, \hat{w}_i, \hat{h}_i)$ and the four parameters denote the centre coordinates, width and height of the predicted bounding box of g_i , respectively. $(\hat{p}_1, \hat{p}_2, \dots, \hat{p}_N)$ can be viewed as a placing trajectory of the character sequence, so we employ an RNN-based **Sequence Discriminator** to distinguish between $\hat{\mathbf{p}}$ and \mathbf{p} under the condition of f_N^c (the final state of \mathbf{f}^c). By performing **Differentiable Composition** (denoted as F_C), we have the synthesized text logo image $\hat{l} = F_C(\mathbf{g}, \hat{\mathbf{p}})$. To capture finer details in the logo images and further ensure that the generated images follow the fundamental principles of human design (*e.g.*, there should be no large-portion overlap between characters in text logos), an **Image Discriminator** is introduced to distinguish between \hat{l} and l under the condition of f_N^c .

3.3. Encoding Visual and Linguistic Information

We take both the visual and linguistic information of inputs into consideration to predict the geometric parameters of each glyph. First, we utilize a CNN network as the visual encoder to extract the visual feature of each glyph image, denoted as $\mathbf{f}^v = (f_1^v, f_2^v, \dots, f_N^v)$. Then, we concatenate the visual feature and the character embedding of each input element to compose a new sequence \mathbf{f}' as $([f_1^v, f_1^e], [f_2^v, f_2^e], \dots, [f_N^v, f_N^e])$, where the square bracket denotes concatenation. Finally, \mathbf{f}' is sent into a condition encoder to obtain the condition features, denoted as $\mathbf{f}^c = (f_1^c, f_2^c, \dots, f_N^c)$. The state at the last step f_N^c can be viewed as a holistic conditional representation of the inputs (images and texts), which is utilized as the condition of dual

discriminators, including both the sequence discriminator and the image discriminator. The whole encoding process can be defined as $\mathbf{f}^c = E(\mathbf{f}^v, \mathbf{f}^e)$. Afterwards, we fed the condition features \mathbf{f}^c and a noise z to the coordinate generator G , to derive the geometric parameters $\hat{\mathbf{p}}$ as $G(\mathbf{f}^c, z)$, where $G(\cdot, \cdot)$ is the coordinate generator approximated by an RNN encoder-decoder; z is randomly sampled from the Standard Normal Distribution and set as the initial state of G 's encoder; \mathbf{f}^c is the input sequence for G . The function of G 's encoder is to propagate the layout style information z to each position of the sequence. The last layer of the coordinate generator is appended with a Sigmoid function whose output is multiplied by W_c or H_c , to guarantee that the predicted parameters are safely in the range of $(0, W_c)$ for \hat{x}_i^c and \hat{w}_i , and $(0, H_c)$ for \hat{y}_i^c and \hat{h}_i .

3.4. Differentiable Composition

After obtaining the predicted geometric parameters corresponding to the glyph images, it is required to transform the glyph images into a text logo with the layout according to the glyphs' geometric parameters. More importantly, the transformation process should not destruct the end-to-end optimization process, *i.e.*, the transformation should propagate gradients. To this end, we render the glyph images on the canvas according to the generated coordinates in a differentiable way on the basis of a variant of STN [10]. To ease the illustration, we first briefly introduce the vanilla STN. Given the source coordinates (x_j^s, y_j^s) of grids in the input feature map, STN aims to learn a set of affine transformation parameters Θ and transform (x_j^s, y_j^s) into the target coordinates (x_j^t, y_j^t) in the output feature maps that are expected to facilitate the down-stream tasks. The transformation is formulated as

$$\begin{pmatrix} x_j^t \\ y_j^t \\ 1 \end{pmatrix} = \Theta \begin{pmatrix} x_j^s \\ y_j^s \\ 1 \end{pmatrix} = \begin{bmatrix} \Theta_{11} & \Theta_{12} & \Theta_{13} \\ \Theta_{21} & \Theta_{22} & \Theta_{23} \end{bmatrix} \begin{pmatrix} x_j^s \\ y_j^s \\ 1 \end{pmatrix}, \quad (1)$$

where j is the index of a grid in the feature maps. With the established mapping, STN performs Differentiable Image Sampling (DIS) implemented by the bilinear interpolation to obtain the transformed feature maps.

Our task is different from the vanilla STN in that we have already had the predicted coordinates $\hat{p}_i = (\hat{x}_i^c, \hat{y}_i^c, \hat{w}_i, \hat{h}_i)$ of target glyphs. We need to obtain the affine transformation parameters and then perform DIS to derive the final rendered glyphs on the canvas. In our scenario, only translation and scaling are considered, so we set $\Theta_{12} = 0$ and $\Theta_{21} = 0$. We use the coordinates of the four corners of the input glyph images $((0, 0), (W_g, 0), (W_g, H_g), (0, H_g))$ and the coordinates of the corresponding transformed images $((\hat{x}_i^c - \frac{\hat{w}_i}{2}, \hat{y}_i^c - \frac{\hat{h}_i}{2}), (\hat{x}_i^c + \frac{\hat{w}_i}{2}, \hat{y}_i^c - \frac{\hat{h}_i}{2}), (\hat{x}_i^c + \frac{\hat{w}_i}{2}, \hat{y}_i^c + \frac{\hat{h}_i}{2}), (\hat{x}_i^c - \frac{\hat{w}_i}{2}, \hat{y}_i^c + \frac{\hat{h}_i}{2}))$ to construct the equation set. Intuitively,

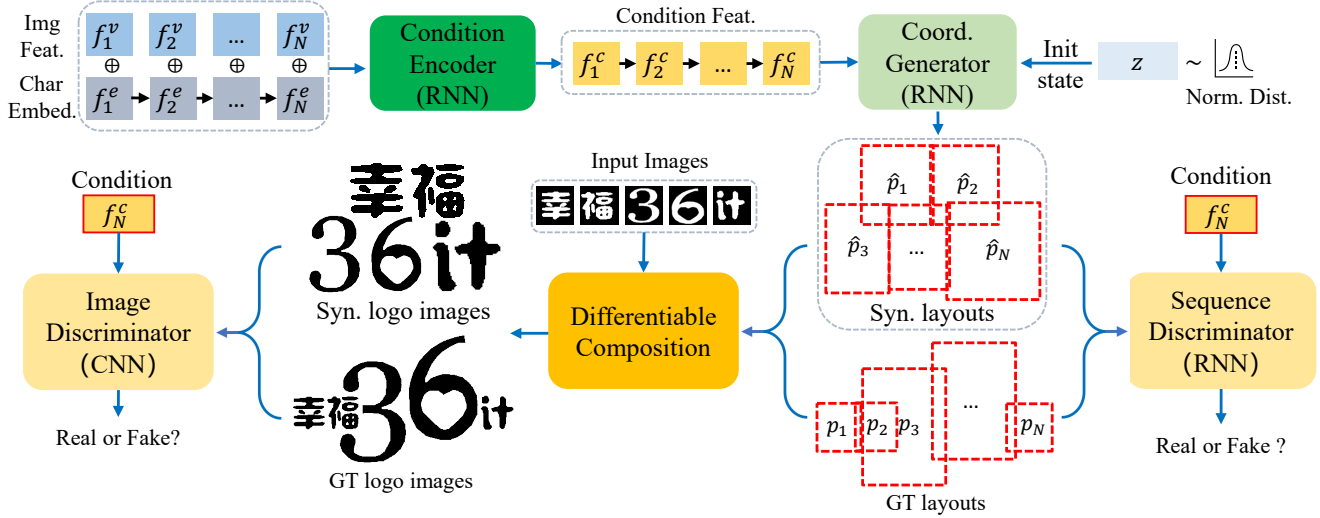


Figure 3. The pipeline of our model. “Char Embed.” denotes char embeddings, “Img Feat.” denotes image features, “Syn.” denotes synthesized, “GT” denotes ground-truth, \oplus denotes the concatenation operation, and the blue arrow represents the direction of data flow.

the equation set can be theoretically resolved, and the transformation parameters are derived as

$$\Theta_i = \begin{bmatrix} \frac{2 \cdot W_g}{\hat{w}_i} & 0 & \frac{(\hat{w}_i - 2 \cdot \hat{x}_i^c) W_g}{\hat{w}_i} \\ 0 & \frac{2 \cdot H_g}{\hat{h}_i} & \frac{(\hat{h}_i - 2 \cdot \hat{y}_i^c) H_g}{\hat{h}_i} \end{bmatrix}. \quad (2)$$

Having the transformation parameters Θ , we apply the corresponding transformation on each glyph image as

$$g'_i = F_A(g_i, \Theta_i), \quad (3)$$

where g'_i is the transformed glyph image of g_i and F_A denotes the transformation function for images implemented as DIS. Typically, the shapes of different glyphs have no overlap in text logos so we do not need to take the layering of elements into consideration. Therefore, we can perform simple additions to obtain the final logo image $\hat{l} \in \mathbb{R}^{H_c \times W_c}$ as

$$\hat{l} = \min\left(\sum_{i=1}^N g'_i, v_{max}\right), \quad (4)$$

where H_c and W_c are the height and width of the logo canvas, respectively. The output pixel values greater than v_{max} are truncated, to guarantee that the pixel values of \hat{l} are in the range of $[0, v_{max}]$ (typically $v_{max} = 255$). An illustration of transformed glyph images and the obtained logo image is shown in Figure 4.

3.5. The Dual Discriminators

Character placing trajectories are supposed to conform with the reading order and possess diverse styles. However, these characteristics cannot be easily captured by CNNs that are commonly used as the discriminator in image synthesis

models. To tackle this, a dual-discriminator module including a Sequence Discriminator and an Image Discriminator is proposed.

The Sequence Discriminator D_s takes the sequence of geometric parameters (\mathbf{p} or $\hat{\mathbf{p}}$) as inputs and the encoded condition f_N^c as the initial state, to analyze the reasonableness of the character placing trajectory. The outputs are denoted as $D_s(\hat{\mathbf{p}}, f_N^c)$ and $D_s(\mathbf{p}, f_N^c)$, representing the probabilities of $\hat{\mathbf{p}}$ and \mathbf{p} being real, respectively.

The Sequence Discriminator is unable to capture fine-grained information (such as strokes) because they receive only the geometric parameters as inputs. Thus, the Image Discriminator D_i is introduced to investigate the details of logo images (synthesized or human-designed) and then predict whether they are real or fake. Following the suggestion of [25], we tile the condition feature f_N^c after the first convolution layer of D_i . The outputs of D_i are denoted as $D_i(\hat{l}, f_N^c)$ and $D_i(l, f_N^c)$, representing the probabilities of \hat{l} and l being real, respectively.

3.6. The Overlap Loss

Generally, the overlap of different glyphs will affect the readability of the synthesized logos and make them look uncoordinated. Through our experiments, we discover that small-portion overlaps cannot be distinguished by the Image Discriminator. Therefore, an overlap loss \mathcal{L}_{ol} is introduced to explicitly penalize the overlap of the transformed glyphs, which is formulated as

$$\mathcal{L}_{ol} = \sum_{i=1}^N (g'_i \otimes (g'_0 \oplus \dots \oplus g'_{i-1})), \quad (5)$$

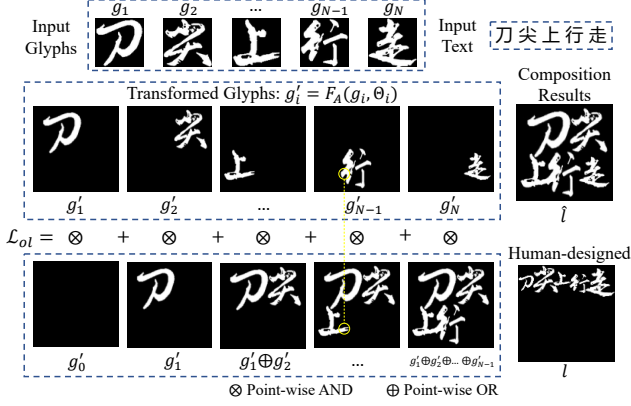


Figure 4. An illustration of transforming elemental glyph images into the canvas and how the overlap loss is calculated. The yellow circles highlight the overlap of different glyphs.

where \otimes denotes the point-wise AND operation and \oplus denotes the point-wise OR operation; g'_0 is a blank mask which can be omitted. An example to illustrate the overlap loss is shown in Figure 4, where the input text is “刀尖上行走”. The transformed glyph g'_4 (“行”) overlaps with the previously rendered glyphs ($g'_1 \oplus g'_2 \oplus g'_3$), which will activate the penalty function \mathcal{L}_{ol} .

3.7. Loss Function

The loss functions for the Sequence and Image Discriminators are:

$$\begin{aligned} \mathcal{L}_D^s &= \log(D_s(\mathbf{p}, f_N^c)) + \log(1 - D_s(\hat{\mathbf{p}}, f_N^c)) \\ &= \log(D_s(\mathbf{p}|E)) + \log(1 - D_s(G(E, z)|E)), \end{aligned} \quad (6)$$

$$\begin{aligned} \mathcal{L}_D^i &= \log(D_i(l, f_N^c)) + \log(1 - D_i(\hat{l}, f_N^c)) \\ &= \log(D_i(l|E)) + \log(1 - D_i(F_C(\mathbf{g}, G)|E)), \end{aligned} \quad (7)$$

respectively, where we omit the inputs of some functions (such as E and G) for brevity. The overall objective function is

$$\min_{E, G} \max_{D_s, D_i} (\mathcal{L}_D^s + \mathcal{L}_D^i + \lambda \mathcal{L}_{ol}), \quad (8)$$

where the aim of optimizing the parameters of E and G is to synthesize a realistic sequence $\hat{\mathbf{p}}$ and its corresponding logo image \hat{l} to fool the Sequence and Image Discriminators, respectively, while the parameters of D_s and D_i are optimized to distinguish $\hat{\mathbf{p}}$ and \hat{l} from the real ones. λ is a hyper-parameter and experimentally chosen.

4. Dataset

We construct a text logo dataset named as **TextLogo3K** by collecting data from Tencent Video, one of the leading online video platforms in China. The dataset consists of 3,470 carefully-selected text logo images that are extracted from the posters/covers of the movies, TV series,



Figure 5. The proposed text logo dataset, consisting of 3,470 text logo images with various styles and annotated with pixel-level segmentation, bounding boxes, and character categories.

and comics. We manually annotate the bounding box, pixel-level mask, and category for each character in those text logos. Some examples are shown in Figure 5. For each example, a color sequence is used to distinguish between different glyphs in the pixel-level annotations of text logos. In the bounding box annotations, we only annotate the most centre parts if the glyph has very long strokes. We also annotate the angle if the glyph is rotated or affine-transformed. It is notable that, the rotation or affine transformation are not considered in our model, but it can be applied in these situations with minor modifications. Besides the layout generation and text logo synthesis tasks, the proposed dataset can also benefit other relevant research (text segmentation, artistic text recognition, artistic font generation, etc.) in the community.

To verify the generality of our model in other writing systems, we make use of the TextSeg dataset proposed in [34], which consists of English text images with pixel-level segmentation. For this dataset, we set the design element as an English word.

5. Experiments

5.1. Implementing Details

The height H_g and width W_g of input glyph images are set to 64. The height H_c and width W_c of the logo canvas are set to 128. The maximum number of N is 20. We split our TextLogo3K dataset into two parts (90% for training and 10% for testing). The visual encoder is a pretrained VGG-19 [30] network which is fine-tuned during training.

5.2. Evaluation Metrics

We adopt two commonly-used metrics for image generation to evaluate the performance of our model and others: Fréchet Inception Distance (FID) [7] and Inception Score (IS) [29]. The logo images rendered with our synthesized layouts (viewed as domain A) and the ground-truth logo images from the whole dataset (viewed as domain B) are compared to calculate FID and IS.



Figure 8. Testing our model’s awareness of both visual and linguistic information. The yellow circles highlight the overlap between different glyphs. ‘LO’ denotes layout. ‘/’ is the symbol for splitting tokens.

Table 2. The effect of the overlap loss. ‘opt.’ denotes optimizing. $\lambda\mathcal{L}_{ol}$ is evaluated on the testing dataset.

Method	$\lambda\mathcal{L}_{ol} \downarrow$	FID \downarrow	IS \uparrow
w/o opt. \mathcal{L}_{ol}	0.135	21.6	2.51
w/ opt. \mathcal{L}_{ol}	0.031	19.7	2.66

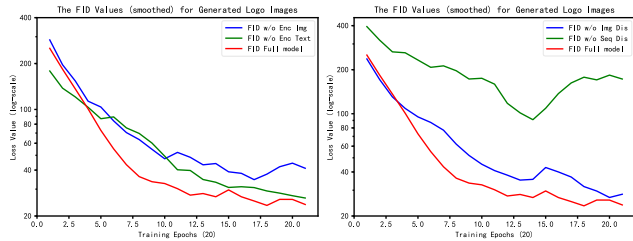


Figure 9. How FID values vary with training epochs under different ablations of our model. ‘Enc Text’ and ‘Enc Img’ denote encoding the input text and images, respectively. ‘Seq Dis’ and ‘Img Dis’ denote the Sequence Discriminator and Image Discriminator, respectively.

input texts and glyph images.

Effect of dual discriminators. As shown in Figure 10, without the Sequence Discriminator, the module tends to generate abnormal sequences that do not conform to human reading habit. For example, the ‘3’ in in ‘叶问3’ moves to the bottom-left of ‘叶’, but the Image Discriminator cannot recognize it as a fake logo image. Without the Image Discriminator, the details (such as spacing and stroke collision) are not as good as the results of the full model. For example, the glyphs in ‘老炮儿’ and ‘黄金时代’ are too close to each other. Quantitative experimental results are shown in Table 1, where removing the Sequence Discriminator (w/o Seq Dis) results in the worst performance.

Effect of the overlap loss. We compute the values of \mathcal{L}_{ol} , FID, and IS on the testing dataset without and with optimizing \mathcal{L}_{ol} when training the model. The results are shown in Table 2, indicating that optimizing \mathcal{L}_{ol} can significantly reduce the overlap and steadily improve FID and IS.

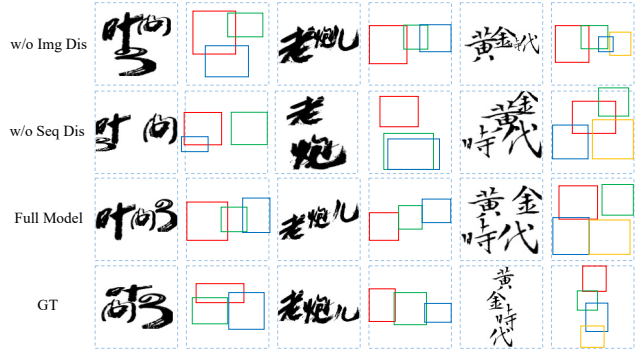


Figure 10. Qualitative results of the ablation studies. ‘Img Dis’ and ‘Seq Dis’ denote the Image and our Sequence Discriminator, respectively. The layout of each logo image is demonstrated by the bounding boxes next to it.

5.6. Comparison Against Other Methods

We compare our method with LayoutGAN [16], LayoutNet [38], and rule-based methods. To make LayoutGAN and LayoutNet applicable in our task, the glyphs are simplified as rectangles and each character category is assigned with a corresponding color. There are three rule-based methods tested in our experiments: (a) arranging all glyphs in a horizontal line; (b) arranging all glyphs in a horizontal or a vertical line (50% probability); (c) arranging glyphs of one or more tokens (separated by Jieba²) in a horizontal/vertical line, and randomly changing the spacing of different lines/glyphs. From Table 1, we can see that our method significantly outperforms others. One major problem of LayoutGAN and LayoutNet is that they cannot generate layouts with correct reading orders, which is not fully reflected in Table 1 but vividly reflected in Figure 11. This is because the image discriminator they both used can only capture the information such as characters’ spacing but not their placing trajectories. They also generate examples with more overlaps than ours, because they do not consider the detailed shapes of input elements. The results also verify that the task of generating aesthetic text logos cannot be properly handled by rule-based methods, whose synthesized layouts are too regular and boring.

5.7. User Study

We conduct a Turing test where the participants (27 professional designers and 52 non-designers) are asked to select one from each pair of logos with the same text but different layouts (our synthesized and human-designed). For each logo pair, they are asked to choose (1) the logo image which they think is generated by machines (AI) (2) the logo image which they think has better quality. There are 20 examples with different layouts and font styles in this test.

²<https://pypi.org/project/jieba/>

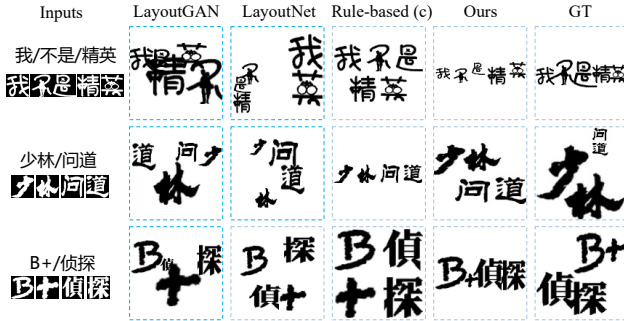


Figure 11. Qualitative comparison among different methods.

Table 3. Results of the user study conducted to evaluate our synthesized layouts. The ratings of “Quality” are from 1 to 5.

User Group	Accuracy↓	Prefer.↑	Quality↑
Designers	58.1%	33.0%	3.0
Non-designers	52.7%	44.9%	3.8
Average	54.6%	40.8%	3.5

After finishing the Turing test, all participants are asked to give ratings (1 to 5, worst to best) for our synthesized layouts in terms of quality. The results are shown in Table 3, where the average accuracy of picking out our synthesized results is 54.6% and the probability of our results being considered better is 40.8%, which indicates that the quality of our synthesized layouts is comparable to human-designed ones.

5.8. Text Logo Synthesis System

Inspired by font generation and texture transfer models ([3, 33, 37], and [21, 22, 36]), we build a text logo synthesis system which takes a desired text and its topic as input and synthesizes the text logo image automatically. Specifically, we first synthesize the font according to the user-input topics, such as lovely, funny, technical, etc. Implementation details about the font synthesis network are described in the supplementary file. Then, we send the synthesized glyph images into our layout generation network to synthesize the text logo image. Finally, we perform the texture transfer method proposed by [22] to obtain the final text logo with texture effects. Several examples are shown in Figure 12, demonstrating the effectiveness of our system for synthesizing aesthetic text logos.

5.9. Limitations

Some failure cases of our method are shown in Figure 13. The proposed layout generation network is not robust enough to glyphs with novel shapes, which could affect our model’s understanding of the text semantics since the image features are concatenated with the character embeddings. Besides, currently our model is still unable to

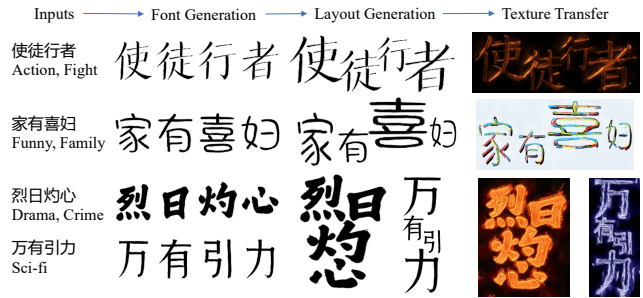


Figure 12. Synthesizing text logos by combining our layout generation method with font generation and texture transfer approaches.

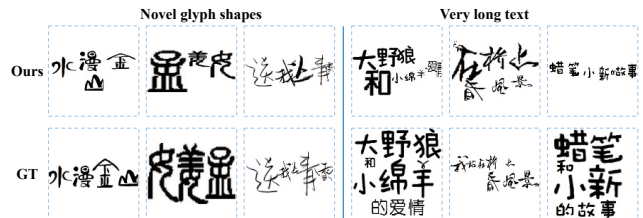


Figure 13. Some failures cases of our model.

satisfactorily handle texts with too many characters/words ($N > 8$), failing to arrange appropriate positions for some glyphs.

6. Conclusion

In this paper, we proposed a layout generation network for aesthetic text logo synthesis. Both the linguistic and visual information of input texts and glyph images were taken into account for layout prediction. Moreover, a novel dual-discriminator module was designed to capture both the placing trajectories and detailed shapes of characters for synthesizing high-quality layouts. We built a large-scale dataset, TextLogo3K, on which extensive experiments (both quantitative and qualitative) were conducted to verify the effectiveness of our method. Finally, a prototype system was developed to automatically synthesize aesthetic text logos by combining our layout generation network with existing font generation and texture transfer models.

Acknowledgements

This work was supported by Beijing Nova Program of Science and Technology (Grant No.: Z191100001119077), Project 2020BD020 supported by PKU-Baidu Fund, Center For Chinese Font Design and Research, Key Laboratory of Science, Technology and Standard in Press Industry, State Key Laboratory of Media Convergence Production Technology and Systems.

References

- [1] Diego Martin Arroyo, Janis Postels, and Federico Tombari. Variational transformer networks for layout generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13642–13652, 2021. **1**
- [2] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. **2**
- [3] Yue Gao, Yuan Guo, Zhouhui Lian, Yingmin Tang, and Jianguo Xiao. Artistic glyph image synthesis via one-stage few-shot learning. *ACM Transactions on Graphics (TOG)*, 38(6):1–12, 2019. **8**
- [4] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. **1, 2**
- [5] Shunan Guo, Zhuochen Jin, Fuling Sun, Jingwen Li, Zhaorui Li, Yang Shi, and Nan Cao. Vinci: An intelligent graphic design system for generating advertising posters. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–17, 2021. **1, 2**
- [6] Kamal Gupta, Justin Lazarow, Alessandro Achille, Larry S. Davis, Vijay Mahadevan, and Abhinav Shrivastava. Layout-transformer: Layout generation and completion with self-attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1004–1014, October 2021. **1**
- [7] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. **5**
- [8] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. **2**
- [9] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. **2**
- [10] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28:2017–2025, 2015. **2, 3**
- [11] Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. Layoutvae: Stochastic scene layout generation from a label set. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9895–9904, 2019. **1**
- [12] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. **1**
- [13] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiri Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. **2**
- [14] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. **2**
- [15] Hsin-Ying Lee, Lu Jiang, Irfan Essa, Phuong B Le, Haifeng Gong, Ming-Hsuan Yang, and Weilong Yang. Neural design network: Graphic layout generation with constraints. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 491–506. Springer, 2020. **1, 2**
- [16] Jianan Li, Jimei Yang, Aaron Hertzmann, Jianming Zhang, and Tingfa Xu. Layoutgan: Generating graphic layouts with wireframe discriminators. *arXiv preprint arXiv:1901.06767*, 2019. **1, 2, 6, 7**
- [17] Jianan Li, Jimei Yang, Jianming Zhang, Chang Liu, Christina Wang, and Tingfa Xu. Attribute-conditioned layout gan for automatic graphic design. *IEEE Transactions on Visualization and Computer Graphics*, 27(10):4039–4048, 2020. **1, 2**
- [18] Runde Li, Jinshan Pan, Zechao Li, and Jinhui Tang. Single image dehazing via conditional generative adversarial network. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. **2**
- [19] Shen Li, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xi-aoyong Du. Analogical reasoning on chinese morphological and semantic relations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 138–143, 2018. **2**
- [20] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P Xing. Dual motion gan for future-flow embedded video prediction. In *The International Conference on Computer Vision (ICCV)*, 2017. **2**
- [21] Yifang Men, Zhouhui Lian, Yingmin Tang, and Jianguo Xiao. A common framework for interactive texture transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6353–6362, 2018. **8**
- [22] Yifang Men, Zhouhui Lian, Yingmin Tang, and Jianguo Xiao. Dyntypo: Example-based dynamic text effects transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5870–5879, 2019. **2, 8**
- [23] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. **2**
- [24] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. **2**
- [25] Guim Perarnau, Joost Van De Weijer, Bogdan Raducanu, and Jose M Álvarez. Invertible conditional gans for image editing. *arXiv preprint arXiv:1611.06355*, 2016. **4**
- [26] Pradyumna Reddy, Paul Guerrero, Matt Fisher, Wilmot Li, and Niloy J Mitra. Discovering pattern structure using differentiable compositing. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020. **2**

- [27] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. In *The International Conference on Learning Representations (ICLR)*, 2016. [2](#)
- [28] Wenqi Ren, Lin Ma, Jiawei Zhang, Jinshan Pan, Xiaochun Cao, Wei Liu, and Ming-Hsuan Yang. Gated fusion network for single image dehazing. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [29] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016. [5](#)
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [5](#)
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [2](#)
- [32] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *The Conference on Neural Information Processing Systems (NeurIPS)*, 2016. [2](#)
- [33] Yizhi Wang, Yue Gao, and Zhouhui Lian. Attribute2font: creating fonts you want from attributes. *ACM Transactions on Graphics (TOG)*, 39(4):69–1, 2020. [2](#), [8](#)
- [34] Xingqian Xu, Zhifei Zhang, Zhaowen Wang, Brian Price, Zhonghao Wang, and Humphrey Shi. Rethinking text segmentation: A novel dataset and a text-specific refinement approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12045–12055, 2021. [5](#)
- [35] Cheng-Fu Yang, Wan-Cyuan Fan, Fu-En Yang, and Yu-Chiang Frank Wang. Layouttransformer: Scene layout generation with conceptual and spatial diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3732–3741, June 2021. [1](#)
- [36] Shuai Yang, Jiaying Liu, Zhouhui Lian, and Zongming Guo. Awesome typography: Statistics-based text effects transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7464–7473, 2017. [8](#)
- [37] Yexun Zhang, Ya Zhang, and Wenbin Cai. Separating style and content for generalized style transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8447–8455, 2018. [8](#)
- [38] Xinru Zheng, Xiaotian Qiao, Ying Cao, and Rynson WH Lau. Content-aware generative modeling of graphic design layouts. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019. [1](#), [2](#), [6](#), [7](#)