

DeepVecFont-v2: Exploiting Transformers to Synthesize Vector Fonts with Higher Quality

Yuqing Wang^{1,2}, Yizhi Wang¹, Longhui Yu², Yuesheng Zhu², Zhouhui Lian^{1*}

¹Wangxuan Institute of Computer Technology, Peking University, China

²School of Electronic and Computer Engineering, Peking University, China

{wyq, yulonghui}@stu.pku.edu.cn, {wangyizhi, zhuys, lianzhouhui}@pku.edu.cn,

Abstract

Vector font synthesis is a challenging and ongoing problem in the fields of Computer Vision and Computer Graphics. The recently-proposed DeepVecFont [27] achieved state-of-the-art performance by exploiting information of both the image and sequence modalities of vector fonts. However, it has limited capability for handling long sequence data and heavily relies on an image-guided outline refinement post-processing. Thus, vector glyphs synthesized by DeepVecFont still often contain some distortions and artifacts and cannot rival human-designed results. To address the above problems, this paper proposes an enhanced version of DeepVecFont mainly by making the following three novel technical contributions. First, we adopt Transformers instead of RNNs to process sequential data and design a relaxation representation for vector outlines, markedly improving the model's capability and stability of synthesizing long and complex outlines. Second, we propose to sample auxiliary points in addition to control points to precisely align the generated and target Bézier curves or lines. Finally, to alleviate error accumulation in the sequential generation process, we develop a context-based self-refinement module based on another Transformer-based decoder to remove artifacts in the initially synthesized glyphs. Both qualitative and quantitative results demonstrate that the proposed method effectively resolves those intrinsic problems of the original DeepVecFont and outperforms existing approaches in generating English and Chinese vector fonts with complicated structures and diverse styles.

*Zhouhui Lian is the corresponding author.

This work was supported by National Language Committee of China (Grant No.: ZDI135-130), Project 2020BD020 supported by PKU-Baidu Fund, Center For Chinese Font Design and Research, and Key Laboratory of Science, Technology and Standard in Press Industry (Key Laboratory of Intelligent Press Media Technology).




Figure 1. Visualization of the vector glyphs synthesized by DeepVecFont and Ours, where different colors denote different drawing commands. (a) DeepVecFont w/o refinement suffers from location shift. (b) DeepVecFont w/ refinement has both over-smoothness (see green circles) and under-smoothness (see blue circles). (c) Our method can directly synthesize visually-pleasing results with compact and coordinated outlines. Zoom in for better inspection.

1. Introduction

Vector fonts, in the format of Scalable Vector Graphics (SVGs), are widely used in displaying documents, arts, and media contents. However, designing high-quality vector fonts is time-consuming and costly, requiring extensive experience and professional skills from designers. Automatic font generation aims to simplify and facilitate the font designing process: learning font styles from a small set of user-provided glyphs and then generating the complete font library. Until now, there still exist enormous challenges due to the variety of topology structures, sequential lengths, and styles, especially for some writing systems such as Chinese.

Recent years have witnessed significant progress [5, 17] made by deep learning-based methods for vector font generation. Nevertheless, vector fonts synthesized by these existing approaches often contain severe distortions and are typically far from satisfactory. More recently, Wang

and Lian [27] proposed DeepVecFont that utilizes a dual-modality learning architecture by exploiting the features of both raster images and vector outlines to synthesize visually-pleasing vector glyphs and achieve state-of-the-art performance. However, DeepVecFont tends to bring location shift to the raw vector outputs (Fig. 1(a)), which are then further refined according to the synthesized images. Specifically, DeepVecFont adopts a differentiable rasterizer [14] to fine-tune the coordinates of the raw vector glyphs by aligning their rasterized results and the synthesized images. However, after the above process, the refined vector outputs tend to be over-fitted to the inherent noise in the synthesized images. Thus, there often exist suboptimal outlines (Fig. 1(b)) with over-smoothed corners (green circles) or under-smoothed adjacent connections (blue circles) in the final synthesized vector fonts, making them unsuited to be directly used in real applications.

To address the above-mentioned issues, we propose a new Transformer-based [25] encoder-decoder architecture, named DeepVecFont-v2, to generate high-fidelity vector glyphs with compact and coordinated outlines. Firstly, we observed that the commonly used SVG representation, which shares the same starting and ending points between adjacent drawing commands, is more suitable for the learning process of RNNs [9] than Transformers [25]. RNNs simulate a recurrent drawing process, where the next movement is determined according to the current hidden state fed with the current drawing command. Therefore, the starting point of the following drawing command can be omitted (replaced by the ending point of the current drawing command). On the contrary, Transformers make drawing prediction based on the self-attention operations performed on any two drawing commands, whether adjacent or not. Therefore, to make the attention operator receive the complete information of their positions, the starting point of each drawing command cannot be replaced by the ending point of the previous command. Based on the above observation, we propose a relaxation representation that models these two points separately and merges them via an extra constraint. Secondly, although the control points of a Bézier curve contain all the primitives, we found that the neural networks still need more sampling points from the curve to perform a better data alignment. Therefore, we sample auxiliary points distributed along the Bézier curves when computing the proposed Bézier curve alignment loss. Thirdly, to alleviate the error accumulation in the sequential generation process, we design a self-refinement module that utilizes the context information to further remove artifacts in the initially synthesized results. Experiments conducted on both English and Chinese font datasets demonstrate the superiority of our method in generating complicated and diverse vector fonts and its capacity for synthesizing longer sequences compared to existing approaches. To summarize,

the major contributions of this paper are as follows:

- We develop a Transformer-based generative model, accompanied by a relaxation representation of vector outlines, to synthesize high-quality vector fonts with compact and coordinated outlines.
- We propose to sample auxiliary points in addition to control points to precisely align the generated and target outlines, and design a context-based self-refinement module to fully utilize the context information to further remove artifacts.
- Extensive experiments have been conducted to verify that state-of-the-art performance can be achieved by our method in both English and Chinese vector font generation.

2. Related Work

Font generation methods can be roughly classified into approaches that aim to synthesize fonts consisting of raster images and vector glyphs, respectively.

In recent years, glyph image synthesis methods typically draw inspirations from recent advances in deep generative models such as VAEs [12] and GANs [7]. Tian et al. [24] and Lyu et al. [18] employed the framework of pix2pix [11] based on cGAN [19] to transfer a template font to target fonts with desired styles. Zhang et al. [29] proposed EMD to explicitly separate the style and content features of glyph images. Mc-GAN [2] and AGIS-Net [6] decomposed the pipeline of artistic font generation into glyph image synthesis and texture transfer. Wang et al. [26] presented Attribute2font, a cGAN-based network to synthesize fonts according to user-specified attributes and their corresponding values. For Chinese font generation, Park et al. [20] and Kong et al. [13] proposed to fully exploit the component and layout information in Chinese glyphs. Xie et al. [28] proposed a deformable generative network which can synthesize target glyph images without direct supervision. Tang et al. [23] adopted a cross-attention mechanism for a fine-grained local style representation to generate target glyph images.

There is growing interest in the task of vector font synthesis, which can directly deliver results with scale-invariant representation. Suveeranont and Igarash [22] proposed to represent the user-defined character example as a weighted sum of the outlines and skeletons from the template fonts, and apply the weights to all characters to generate the new font. Campbell and Kautz [4] aimed to learn a font manifold from existing fonts and create new font styles by interpolation and extrapolation from the manifold. SketchRNN [8] employed a sequence VAE based on bi-directional RNNs to generate sketches. Easyfont [15] decomposed the handwriting style into the shape and layout style of strokes to

generate personal handwriting Chinese characters. SVG-VAE [17] developed an image autoencoder architecture to learn style vectors of fonts, and then used LSTMs [9] followed by a Mixture Density Network [3] to generate the SVG drawing sequence. DeepSVG [5] adopted a hierarchical generative network based on Transformers to generate vector icons with multiple paths. Im2Vec [21] can directly generate vector graphics from raster training images without explicit vector supervision. DeepVecFont [27] exploited both the image and sequence modalities to facilitate the synthesis of vector glyphs. Liu et al. [16] proposed to represent glyphs implicitly as shape primitives enclosed by several quadratic curves, followed by a style transfer network to render glyph images at arbitrary resolutions. More recently, Aoki and Aizawa [1] introduced AdaIn [10] into the pipeline of DeepSVG to synthesize Chinese vector fonts. However, there still often exist non-negligible distortions and artifacts on the vector glyphs synthesized by the above-mentioned existing approaches.

3. Method

In this section, we present the details of our proposed DeepVecFont-v2. Specifically, we first introduce the data structure and the learned SVG embeddings of vector glyphs. Then, we describe the overview of our network architecture and the implementation details of each module.

3.1. Data Structure and SVG Embeddings

A vector font is a set of vector glyphs $\{G_1, \dots, G_{N_{char}}\}$, where N_{char} denotes the number of character categories. As shown in Fig. 3, a vector glyph G_i can be viewed as a sequence of drawing commands, denoted as $G_i = [C_{i,1}, \dots, C_{i,N_c}]$, where N_c is the total length of commands. The drawing command $C_{i,j}$ is a tuple $(z_{i,j}, p_{i,j})$, where $z_{i,j}$ and $p_{i,j}$ denote the command type and command coordinates, respectively. The drawing sequence starts from the top-left command. We consider 4 command types, i.e., $z_i \in \{MoveFromTo, LineFromTo, CurveFromTo, EOS\}$. *MoveFromTo* means moving the drawing position to a new location, which is used for starting a new path. *LineFromTo* and *CurveFromTo* mean drawing a straight line and a three-order Bézier curve, respectively. *EOS* means ending the drawing sequence.

Typically, $p_{i,j}$ is made up of N_p pairs of coordinates: $p_{i,j} = [(x_{i,j}^1, y_{i,j}^1), \dots, (x_{i,j}^{N_p}, y_{i,j}^{N_p})]$, known as the control points, where N_p is determined by the order of Bézier curves. In the typical SVG representation, the starting point is omitted so that the number of curve order is equal to N_p , namely, $N_p = 3$ for the three-order Bézier curve and $N_p = 1$ for the line (equal to the one-order Bézier curve).

Relaxation Representation Different from SVG-VAE and DeepVecFont, we assign each drawing command

with individual starting and ending points. Therefore, the number of coordinate pairs N_p in *CurveFromTo* is set to 4. Specifically, $p_{i,j}^1$ and $p_{i,j}^4$ are the starting and ending points, respectively; $p_{i,j}^2$ and $p_{i,j}^3$ are two intermediate control points. We pad the length of other commands to that of *CurveFromTo*: for $z_{i,j} \in \{MoveFromTo, LineFromTo\}$, only $(x_{i,j}^1, y_{i,j}^1)$ and $(x_{i,j}^4, y_{i,j}^4)$ (starting and ending points) are used; for $z_{i,j} = EOS$, no argument is used. Afterwards, we render those vector glyphs to obtain the rasterized glyph images.

An illustration of the difference between an existing SVG representation and ours can be found in Fig. 4. Our proposed representation describes each drawing command separately, which is well-suited for the Transformer’s attention mechanism to exploit the long-range dependencies of drawing commands. In the training stage, we force the ending point of C_i to be consistent with the starting point of C_{i+1} using an extra constraint calculated by Eq. 10. In the inference stage, we merge these two points by averaging their positions.

Embedding We first project the drawing command into a common continuous d_E -dimensional embedding space. Specifically, each $C_{i,j}$ is projected into a vector $e_{i,j} \in \mathbb{R}^{d_E}$ via the sum of four embeddings:

$$e_{i,j} = e_{i,j}^{cmd} + e_{i,j}^{args} + e_{i,j}^{w,h} + e_{i,j}^{pos}. \quad (1)$$

For the command type embedding $e_{i,j}^{cmd}$, we use a learnable matrix $W_{cmd} \in \mathbb{R}^{d_E \times 4}$ to convert the command type into a d_E -dimensional vector, which is formulated as $e_{i,j}^{cmd} = W_{cmd} \delta_{i,j}^c \in \mathbb{R}^{d_E}$, where $\delta_{i,j}^c \in \mathbb{R}^4$ is a one-hot vector for the four command types.

For the argument (coordinate) embedding $e_{i,j}^{args}$, we first quantize the continuous coordinates into discrete integers and convert the integer into a one-hot vector with 256 dimensions. Then, we stack all the 8 coordinate parameters into a matrix $\delta_{i,j}^p \in \mathbb{R}^{256 \times 8}$, and embed each parameter using a learnable matrix $W_{args}^b \in \mathbb{R}^{d_E \times 256}$. After that, we aggregate all the parameter embeddings through a linear projection layer $W_{args}^a \in \mathbb{R}^{d_E \times 8d_E}$, which is formulated as:

$$e_{i,j}^{args} = W_{args}^a \text{flatten}(W_{args}^b \delta_{i,j}^p), \quad (2)$$

where $\text{flatten}(\cdot)$ means flattening the input into a vector.

The third term $e_{i,j}^{w,h}$ encodes the height and width of the glyph area to capture global styles. We discretize the height and width into discrete integers and project them into the continuous space to obtain $e_{i,j}^w$ and $e_{i,j}^h$. Then we concatenate them by element-wise addition to get $e_{i,j}^{w,h}$.

The fourth term (positional embedding) $e_{i,j}^{pos}$ encodes the position and order information of all commands in a drawing sequence. Similar to [25], we use the absolute positional encoding to compute the $e_{i,j}^{pos} \in \mathbb{R}^{d_E}$ for each command.

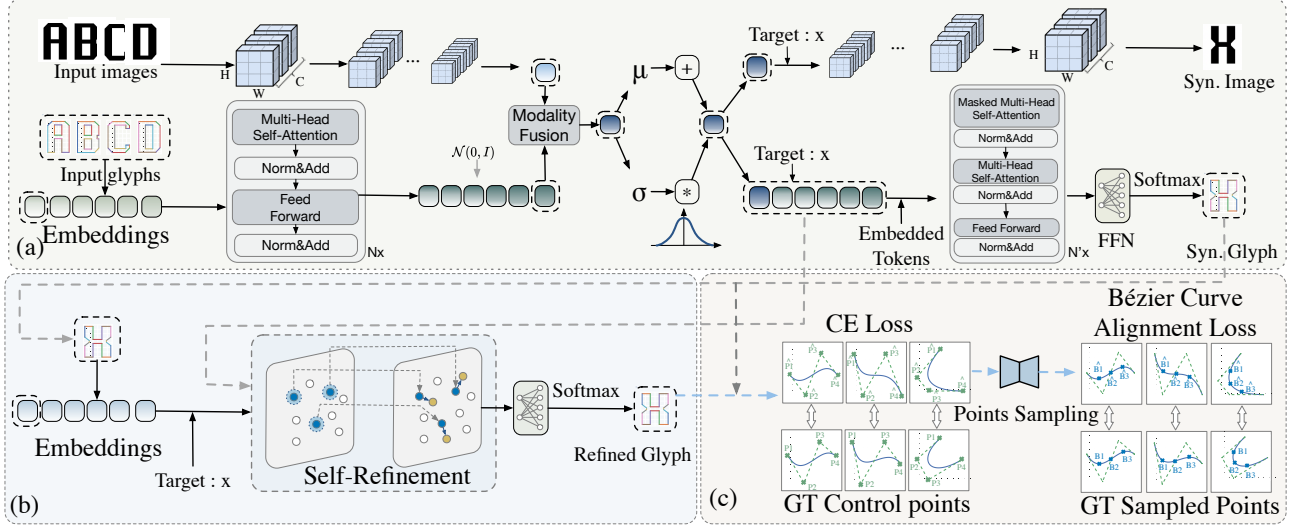


Figure 2. The pipeline of our DeepVecFont-v2. The inputs are reference glyphs in both raster images and vector outlines. (a) A dual-branch architecture based on Transformers and CNNs aims to synthesize the target vector glyph. (b) The self-refinement module is designed to remove artifacts in the initially synthesized vector glyphs. (c) In addition to control points, auxiliary points are sampled to align the synthesized glyph with the corresponding target via the Bézier curve alignment loss.




Figure 3. An illustration of our data structure used to describe vector glyphs. The command type and coordinates are shown below each canvas. The origin of the axes is at the top left of each canvas and “ \emptyset ” denotes the unused argument.

3.2. Dual-branch Pipeline

Fig. 2 shows the pipeline of our DeepVecFont-v2. Given randomly sampled N_r (typically set to 4) reference glyphs as input, the model generates a target glyph with the same font style as input samples, which is further refined by a self-refinement module. Similar to DeepVecFont [27], the proposed model also adopts a dual-branch architecture which, however, employs several new techniques and modules. More details are presented as follows:

Encoder The encoder of our DeepVecFont-v2 is made




Figure 4. A demonstration of different SVG representations. (a) The ground truth. If the previous blue segment is wrongly predicted, (b) the commonly used representation in SVG-VAE [17] shares the connected points (E1 and S2), resulting in location shift. (c) The proposed relaxation representation models these two points separately, which is more robust against outliers.

up with an image encoder and a sequence encoder. The image encoder is a CNN, outputting the image feature $f_{img} \in \mathbb{R}^{d_E}$. Different from DeepVecFont, the sequence encoder is a Transformer encoder composed of six layers of transformer encoder blocks. For each vector glyph G_i , it receives as input the sequence embedding $[e_{i,0}, e_{i,1}, \dots, e_{i,N_c}]$, where $[e_{i,1}, \dots, e_{i,N_c}]$ is the embedding of N_c drawing commands and $e_{i,0}$ denotes an auxiliary learnable “token” for performing dual-modality fusion. The output of our Transformer encoder is denoted as $e'_i = [e'_{i,0}, e'_{i,1}, \dots, e'_{i,N_c}] \in \mathbb{R}^{d_E \times (N_c+1)}$. Next, we calculate the holistic sequence-aspect style feature $f_{seq} = [f_0^{seq}, f_1^{seq}, \dots, f_{N_c}^{seq}]$, where f_j^{seq} is aggregated from all the j -th tokens in e' of N_r reference glyphs via linear projection.

Modality fusion We use the modality token f_0^{seq} to represent the style feature of vector glyphs, and combine it with the image feature f_{img} by a linear projection layer to get the fused feature f :

$$f = Linear([f^{img}; f_0^{seq}]), \quad (3)$$

then f is normalized by using the reparametrization trick introduced in VAE [12].

Decoder The image decoder is a Deconvolutional Neural Network. We send f into the image decoder to generate the target glyph image I_t . Then we employ the L1 loss and perceptual loss to compute the image reconstruction loss:

$$L_{img} = \|\hat{I}_t - I_t\|_1 + L_{percep}(\hat{I}_t, I_t). \quad (4)$$

We feed the sequence decoder with the input of $[f, f_1^{seq}, \dots, f_{N_c}^{seq}]$, where the original modality token f_0^{seq} is replaced with the fused feature f . An MLP layer is appended on top of the Transformer decoder, predicting the command types and coordinates of the target glyph as $[\hat{z}_{t,1}, \dots, \hat{z}_{t,N_c}]$ and $[\hat{p}_{t,1}, \dots, \hat{p}_{t,N_c}]$, respectively. Here we define the loss between the initially generated glyph and its corresponding ground truth as:

$$L_{CE}^{init} = \sum_{j=1}^{N_c} w_{cmd} \ell(z_{t,j}, \hat{z}_{t,j}) + \ell(p_{t,j}, \hat{p}_{t,j}), \quad (5)$$

where ℓ denotes the Cross-Entropy loss, all the coordinates (p) are quantized to be optimized by ℓ , w_{cmd} means the loss weight for command type prediction, and all the unused arguments are masked out in the second item.

3.3. Context-based Self-refinement

In the inference stage, Transformers follow the sequential generation process adopted in RNNs: taking the previous predicted results as the condition to predict the following steps. However, a prediction error generated in a certain step of the sequential process might finally lead to the accumulation of large errors and bring significant distortions in the synthesized glyphs. We notice that there are strong priors and correlations in the geometries of glyphs, such as symmetry and smoothly varied stroke widths, and these context information can be used as strong clues to enhance the glyph synthesizing performance.

Therefore, we propose a self-refinement module to refine the initial predictions by analyzing their context information, i.e., $((\bar{z}_{t,1}, \bar{p}_{t,1}), \dots, (\bar{z}_{t,N_c}, \bar{p}_{t,N_c})) = F_r((\hat{z}_{t,1}, \hat{p}_{t,1}), \dots, (\hat{z}_{t,N_c}, \hat{p}_{t,N_c}))$, where F_r denotes the function of our refinement module which is actually a 2-layer Transformer decoder. The multi-head attention operation (\mathbf{F}_{mha}) used in our decoder layers is formulated as:

$$\mathbf{F}_{mha} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\mathcal{C}(\mathbf{Q})} + \mathbf{M}\right)\mathbf{V}, \quad (6)$$

$$\mathbf{M}_{ij} = \begin{cases} 0, & i \leq \hat{N}_c \\ -\infty, & i > \hat{N}_c \end{cases}, \quad (7)$$

where \hat{N}_c is the length of the predicted sequence; \mathcal{C} is the feature dimension of queries for normalization; $\mathbf{Q} \in \mathbb{R}^{d_E \times N_c}$ consists of the embeddings of the initially predicted sequence $(\hat{z}_{t,1}, \hat{p}_{t,1}), \dots, (\hat{z}_{t,N_c}, \hat{p}_{t,N_c})$; $\mathbf{K}, \mathbf{V} \in \mathbb{R}^{d_E \times (N_c+1)}$ denote the linear projections of memories used in the previous sequence decoder; \mathbf{M} is used to perform self-attention only on the valid positions of the initially predicted sequence. Finally, we refine the command types and coordinates $(\bar{z}_{t,j}, \bar{p}_{t,j})$ again via the supervision of the corresponding ground truth by minimizing:

$$L_{CE}^{refine} = \sum_{j=1}^{N_c} (w_{cmd} \ell(z_{t,j}, \bar{z}_{t,j}) + \ell(p_{t,j}, \bar{p}_{t,j})). \quad (8)$$

3.4. Bézier Curve Alignment Loss

Through our experiments, we found that it is still insufficient by only employing the control points (p) to supervise the alignment of Bézier curves/lines. Therefore, we propose to sample more points from each drawing command $C_{t,j}$. As we know, a Bézier curve built upon control points p is defined as: $B = \sum_{k=0}^n \binom{n}{k} r^k (1-r)^{n-k} p$, where n

denotes the number of curve order, $\binom{n}{k}$ denotes the binomial coefficients, and r is the position parameter. Thus, we can calculate the Bézier curve alignment loss by:

$$L_{bezier} = \sum_{r \in r_p} (\hat{B}_{t,j}(r) - B_{t,j}(r))^2 + (\bar{B}_{t,j}(r) - B_{t,j}(r))^2, \quad (9)$$

where $r_p = \{0.25, 0.5, 0.75\}$ represents the parameters of the auxiliary points we sampled. Note that ‘‘LineFromTo’’ can be formulated as a special case of ‘‘CurveFromTo’’, where all the control points are uniformly distributed on a line segment.

3.5. Merging Relaxed Drawing Commands

Ideally, the starting point of the current drawing command (*LineFromTo* or *CurveFromTo*, l or c for short) should coincide with the previous one’s ending point for both initially synthesized and refined sequences. Thereby, we try to minimize the L2 distance between the predicted

Base model	M	O	I	D	e	c
+ Relaxation Representation	M	O	I	D	e	c
+ Bézier Curve Alignment Loss	M	O	I	D	e	c
+ Self-Refinement	M	O	I	D	e	c

Ground Truth	M	O	I	D	e	c

Figure 5. The ablation study of our method. Blue circles indicate the shortcomings of each incomplete configuration.

ending point of command C_{i-1} and the starting point of C_i :

$$L_{cons} = \mathbf{1}_{z_{t,j} \in \{l,c\}} \sum_{j=2}^{N_c} (\hat{x}_{t,j}^1 - \hat{x}_{t,j-1}^4)^2 + (\hat{y}_{t,j}^1 - \hat{y}_{t,j-1}^4)^2 + (\bar{x}_{t,j}^1 - \bar{x}_{t,j-1}^4)^2 + (\bar{y}_{t,j}^1 - \bar{y}_{t,j-1}^4)^2. \quad (10)$$

3.6. Overall Objective Loss

Combining all losses mentioned above, we train the whole model by minimizing the following objective:

$$L_{img} + L_{CE}^{init} + L_{CE}^{refine} + L_{cons} + L_{bézier} + L_{kl}, \quad (11)$$

where the last term represents the KL loss normalizing the latent code f . For brevity, the weight of each term is omitted from the equation.

4. Experiments

4.1. Dataset and Implementation Details

The dataset used in our experiment for generating English fonts is the same as [27], which includes 8035 fonts for training and 1425 fonts for testing. For Chinese font synthesis, we built a dataset consisting of 212 and 34 fonts for training and testing. To directly apply the same model in two datasets, here we only randomly selected 52 Chinese characters with relatively simple structures. Since the scale of our dataset for Chinese fonts is relatively small, we augmented the training set by applying the affine transformation, enlarging it by ten times.

We conduct the experiments by following the experimental settings of DeepVecFont [27] for fair comparison. Specifically, the numbers of reference glyphs are chosen as 4 and 8 for English and Chinese font generation, respectively. We employ the Adam optimizer with an initial learning rate of 0.0002. The image resolution is chosen as 64×64

in both the training and testing stages. When inferencing, we first add a noise vector distributed by $\mathcal{N}(0, I)$ to the sequence feature, to simulate the feature distortion caused by the human-designing uncertainty (as observed in DeepVecFont). Then, we sample N_s (10 for English and 50 for Chinese) synthesized vector glyphs as candidates and select the one as the final output that has the highest IOU value with the synthesized image. The reconstruction errors (denoted as ‘‘Error’’) in our quantitative results are obtained by computing the average L1 distance between the rasterized image of each synthesized vector glyph and its corresponding ground-truth glyph image (at the resolution of 64×64) in the testing set. In our experiments, the weight values of different losses in Eq. 11 are set to 1.0, 1.0, 1.0, 10, 1.0, and 0.01, respectively from left to right.

4.2. Ablation Study

We conduct qualitative and quantitative experiments to examine the impact of each module in our proposed model. In Tab. 5, the base model is modified from DeepVecFont by simply replacing the LSTM networks with Transformers, then we evaluate each module by adding them successively to the base model. As shown in Fig. 5, the base model using the SVG representation in SVG-VAE [17] tends to generate incomplete glyphs with severe distortions. By employing the relaxation representation, semantic information can be preserved and the smooth connection between adjacent drawing commands is guaranteed. However, there still exist suboptimal curves with slight distortions. After adding the Bézier curve alignment loss to the current model, most distortions on the synthesized curves can be eliminated. Finally, we perform self-refinement to remove artifacts in the initially synthesized glyphs, resulting in high-quality glyphs with compact and coordinated outlines. Tab. 1 shows some quantitative results of our ablation study, further demonstrating the superiority of our proposed modules.

4.3. Parameter Study

We also conduct parameter studies to find the best choice of the number of sampling points distributed along the Bézier curves, and the results are shown in Tab. 3. we can see that the model’s performance will be markedly improved when the number of sampling points changes from 0 to 3, while the performance is just slightly enhanced when we increase the point number from 3 to 9. This is mainly due to the fact that 3 sampling points are sufficient enough to precisely align two Bézier curves/lines. Therefore, we choose to sample 3 points for each Bézier curve/line to achieve a balance between the effectiveness and efficiency of our method. Furthermore, we also conduct parameter studies to examine the performance of our method under different numbers of input reference glyphs. Please refer to the supplemental materials for more details.



Figure 6. Comparison of few-shot font generation results for DeepSVG, DeepVecFont, and our method. Our method shows superiority in generating high-quality English vector fonts with various styles. Please zoom in for better inspection.



Figure 7. Comparison of few-shot Chinese vector font generation results obtained by different methods.

Model	Error-EN↓
Base model	0.0588
+ Relaxation Representation	0.0557
+ Bézier Curve Alignment Loss	0.0529
+ Self-Refinement	0.0519

Table 1. Comparison of reconstruction errors for our methods under different configurations.

4.4. Font Interpolation

Our method can be used to achieve smooth interpolation between two vector fonts. Specifically, given two vector fonts a and b , an interpolated style feature between them can be calculated by:

$$f_{inter} = (1 - \lambda) \cdot f(a) + \lambda \cdot f(b). \quad (12)$$

Then, we feed the interpolated feature into the sequence and image decoders to generate the interpolated vector glyphs and the corresponding glyph images. Fig. 8 shows two font interpolation results for English and Chinese vector fonts, respectively. We can see that glyphs in the source styles can smoothly morph to the target ones, demonstrating the



Figure 8. English and Chinese vector font interpolation results, where the weight, width and styles change smoothly.

Model	Error-EN↓	Error-CN↓
DeepSVG	0.125	0.167
DeepVecFont	0.056	0.086
Ours	0.052	0.080

Table 2. Comparison of reconstruction errors for different methods. “EN” and “CN” denote the English and Chinese testing datasets, respectively.

Point Num	Error-EN↓	Point Num	Error-EN↓
0	0.0557	6	0.0526
1	0.0543	9	0.0524
3	0.0529	12	0.0520

Table 3. Comparison of reconstruction errors under different numbers of sampling points. “EN” denotes the English testing dataset.

effectiveness of our method in vector font interpolation.




Figure 9. Comparison of more vector fonts synthesized by DeepVecFont (DVF) and our DeepVecFont-v2 in the task of few-shot font generation. The input reference glyphs are filled in black. Please zoom in for better inspection.

4.5. Comparison with the State of the Arts.

We compare the performance of our DeepVecFont-v2 with other existing methods (i.e., DeepSVG [5] and DeepVecFont [27]) on English and Chinese vector font datasets. Some qualitative results are shown in Fig. 6, Fig. 7, and Fig. 9. The quantitative results are shown in Tab. 2.

English vector font generation. From Fig. 6, we can see that the vector glyphs synthesized by DeepSVG often contain severe distortions (marked in yellow). This is mainly because the hierarchical architecture of DeepSVG was originally designed for synthesizing vector icons with multiple separated short-length paths, while a vector glyph typically consists of several long-range outlines. The synthesized glyphs of DeepVecFont filled with black pixels generally look visually pleasing. However, there still exist some problems in the synthesis results regarding the details of outlines: 1) the refinement post-processing tends to over-fit the initially synthesized vector glyph with the corresponding raster image, resulting in suboptimal outlines with over-smoothed corners (marked in red); 2) when a sequence with redundant drawing commands is predicted, the outline tends to be staked together with self-interactions (marked in gray); 3) structurally-incorrect glyphs might be synthesized for some special styles (e.g., “F” in the green box). On the contrary, our method can synthesize high-quality vector glyphs for almost all kinds of font styles. Fig. 9 shows more results obtained by our method and DeepVecFont. From Fig. 9, we can also observe that our DeepVecFont-v2 sometimes can even generate vector glyphs whose font styles are more consistent with the input samples compared to the ground truth (marked in blue).

Chinese vector font generation. Chinese glyphs typically contain complex shapes and structures, making the task of Chinese vector font synthesis much more challenging. As shown in Fig. 7, DeepSVG tends to generate glyphs with inconsistent styles and severe artifacts, and

DeepVecFont may obtain fragmented results when synthesizing glyphs with multiple closed paths (marked in green). In contrast, our method synthesizes visually pleasing vector fonts, mainly due to the utilization of Transformers and our specifically-designed modules to guarantee the smoothness of outlines. Fig. 9 shows more results obtained by our method and DeepVecFont. The quantitative results shown in Tab. 2 further demonstrate the superiority of our method when handling fonts consisting of glyphs with more complex topologies and longer drawing-command sequences. More synthesized results can be found in our supplemental materials.

4.6. Limitations

Our model fails to synthesize correct results when handling complex glyphs that contain large amounts of long drawing paths (e.g., “霸” shown in our supplemental materials). One possible reason is that longer sequences inherently exist more uncertainties generated during the human font-designing process [27], bringing more challenges to train the model. Another reason is the possible existence of complicated topological changes for a Chinese character in different styles, making it hard to learn stable SVG embeddings. We leave these issues as our future work.

5. Conclusion

This paper proposed a novel method, DeepVecFont-v2, to effectively handle the challenging task of vector font synthesis. Specifically, we replaced RNNs adopted in the original DeepVecFont by Transformers and adopted several specifically-design modules, including relaxation representation, the Bézier curve alignment loss, and context-based self-refinement. Thus, vector fonts with higher quality can be directly synthesized by our method in an end-to-end manner. Experimental results demonstrated the superiority of our DeepVecFont-v2 compared to the state of the art in the applications of English and Chinese vector font synthesis.

References

- [1] Haruka Aoki and Kiyoharu Aizawa. Svg vector font generation for chinese characters with transformer. *arXiv preprint arXiv:2206.10329*, 2022. 3
- [2] Samaneh Azadi, Matthew Fisher, Vladimir G Kim, Zhaowen Wang, Eli Shechtman, and Trevor Darrell. Multi-content gan for few-shot font style transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7564–7573, 2018. 2
- [3] C. M. Bishop. Mixture density networks. *IEEE Computer Society*, 1994. 3
- [4] Neill DF Campbell and Jan Kautz. Learning a manifold of fonts. *ACM Transactions on Graphics (ToG)*, 33(4):1–11, 2014. 2
- [5] Alexandre Carlier, Martin Danelljan, Alexandre Alahi, and Radu Timofte. Deepsvg: A hierarchical generative network for vector graphics animation. *Advances in Neural Information Processing Systems*, 33:16351–16361, 2020. 1, 3, 8
- [6] 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. 2
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Y. Bengio. Generative adversarial nets. In *Neural Information Processing Systems*, 2014. 2
- [8] David Ha and Douglas Eck. A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477*, 2017. 2
- [9] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 2, 3
- [10] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. 3
- [11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2
- [12] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv.org*, 2014. 2, 5
- [13] Yuxin Kong, Canjie Luo, Weihong Ma, Qiyuan Zhu, Sheng-gao Zhu, Nicholas Yuan, and Lianwen Jin. Look closer to supervise better: One-shot font generation via component-based discriminator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13482–13491, 2022. 2
- [14] T. M. Li, M Luká, M. Gharbi, and J. Ragan-Kelley. Differentiable vector graphics rasterization for editing and learning. *ACM Transactions on Graphics*, 39(6):1–15, 2020. 2
- [15] Zhouhui Lian, Bo Zhao, Xudong Chen, and Jianguo Xiao. Easyfont: a style learning-based system to easily build your large-scale handwriting fonts. *ACM Transactions on Graphics (TOG)*, 38(1):1–18, 2018. 2
- [16] Ying-Tian Liu, Yuan-Chen Guo, Yi-Xiao Li, Chen Wang, and Song-Hai Zhang. Learning implicit glyph shape representation. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 3
- [17] Raphael Gontijo Lopes, David Ha, Douglas Eck, and Jonathon Shlens. A learned representation for scalable vector graphics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7930–7939, 2019. 1, 3, 4, 6
- [18] Pengyuan Lyu, Xiang Bai, Cong Yao, Zhen Zhu, Tengeng Huang, and Wenyu Liu. Auto-encoder guided gan for chinese calligraphy synthesis. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 1095–1100. IEEE, 2017. 2
- [19] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 2
- [20] Song Park, Sanghyuk Chun, Junbum Cha, Bado Lee, and Hyunjung Shim. Few-shot font generation with localized style representations and factorization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2393–2402, 2021. 2
- [21] Pradyumna Reddy, Michael Gharbi, Michal Lukac, and Niloy J Mitra. Im2vec: Synthesizing vector graphics without vector supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7342–7351, 2021. 3
- [22] Rapee Suveeranont and Takeo Igarashi. Example-based automatic font generation. In *International Symposium on Smart Graphics*, pages 127–138. Springer, 2010. 2
- [23] Licheng Tang, Yiyang Cai, Jiaming Liu, Zhibin Hong, Mingming Gong, Minhu Fan, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Few-shot font generation by learning fine-grained local styles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7895–7904, 2022. 2
- [24] Yuchen Tian. zi2zi: Master chinese calligraphy with conditional adversarial networks. *Internet*] <https://github.com/kaonashi-tyc/zi2zi>, 3, 2017. 2
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *arXiv*, 2017. 2, 3
- [26] 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
- [27] Yizhi Wang and Zhouhui Lian. Deepvecfont: synthesizing high-quality vector fonts via dual-modality learning. *ACM Transactions on Graphics (TOG)*, 40(6):1–15, 2021. 1, 2, 3, 4, 6, 8
- [28] Yangchen Xie, Xinyuan Chen, Li Sun, and Yue Lu. Dgfont: Deformable generative networks for unsupervised font generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5130–5140, 2021. 2
- [29] 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. 2