

DeepCalliFont: Few-Shot Chinese Calligraphy Font Synthesis by Integrating Dual-Modality Generative Models

Yitian Liu, Zhouhui Lian*

Wangxuan Institute of Computer Technology, Peking University, Beijing, P.R. China
{lsflyt, lianzhouhui}@pku.edu.cn


Abstract

Few-shot font generation, especially for Chinese calligraphy fonts, is a challenging and ongoing problem. With the help of prior knowledge that is mainly based on glyph consistency assumptions, some recently proposed methods can synthesize high-quality Chinese glyph images. However, glyphs in calligraphy font styles often do not meet these assumptions. To address this problem, we propose a novel model, DeepCalliFont, for few-shot Chinese calligraphy font synthesis by integrating dual-modality generative models. Specifically, the proposed model consists of image synthesis and sequence generation branches, generating consistent results via a dual-modality representation learning strategy. The two modalities (i.e., glyph images and writing sequences) are properly integrated using a feature recombination module and a rasterization loss function. Furthermore, a new pre-training strategy is adopted to improve the performance by exploiting large amounts of uni-modality data. Both qualitative and quantitative experiments have been conducted to demonstrate the superiority of our method to other state-of-the-art approaches in the task of few-shot Chinese calligraphy font synthesis. The source code can be found at <https://github.com/lshlyt-pku/DeepCalliFont>.

Introduction

Manually designing a Chinese font is time-consuming and costly, mainly because there exist huge amounts of Chinese characters. For example, the latest official character set GB18030-2022 contains 87,887 Chinese glyphs, most of which possess delicate shapes and complex structures. In the last decade, large numbers of font synthesis methods (Tian 2017; Jiang et al. 2017; Lian et al. 2018) have been proposed aiming to automatically generate the complete Chinese font library from a small set of glyphs written or designed by a user. However, existing approaches still have the following two main drawbacks: 1) the number of characters required for training is still large; 2) the quality of machine-generated Chinese fonts is still unsatisfactory for many special font styles, especially cursive Calligraphy styles.

As mentioned above, the task of few-shot Chinese font generation is challenging but of great practical value. There



(a) Assumption 1: the same characters in different fonts have the same components and stroke order. (b) Assumption 2: the same components in different characters but the same font should be identical or similar.

Figure 1: Two assumptions of Chinese character prior knowledge. Regular Chinese fonts satisfy these assumptions, but calligraphy fonts don't.

exist some works on this task, like EMD (Zhang, Zhang, and Cai 2018), AGIS-Net (Gao et al. 2019), etc., which concentrate on extracting and integrating features via well-designed loss functions or modules. Meanwhile, there also exist other methods that utilize prior knowledge of Chinese characters to extract structure-aware style representations (Park et al. 2020, 2021; Liu et al. 2022) or encode prior knowledge into content labels (Sun et al. 2017; Wu, Yang, and Hsu 2020; Huang et al. 2020). In this manner, prior knowledge guided methods significantly improve the generation quality of most fonts and keep the fine-grained style consistency. Fig. 1 shows two implicit consistency assumptions of such prior knowledge: 1) the same characters in different font styles have the same components and stroke order; 2) the same components in different characters but the same font style should be identical or similar. However, these methods often fail to satisfactorily handle calligraphy fonts with connected strokes. This is mainly due to the fact that calligraphy font designers tend to change the stroke order or simplify some strokes for better visual aesthetics, leading to calligraphy fonts conflicting with the component and stroke consistency. It can be observed from Fig. 1 that glyphs in calligraphy styles are far more complex and hard to synthesize than those in regular font styles.

Compared to the prior knowledge, a writing trajectory (the sequence of key points) owns closer relationship and better consistency with the corresponding glyph image.

*Corresponding author.

Therefore, we try to use the writing trajectories instead of the prior knowledge and propose a dual-modality few-shot Chinese font generation model, DeepCalliFont, which processes and synthesizes Chinese characters represented in two modalities: glyph images and writing trajectories. More specifically, we map the dual-modality features to the same space using a distillation-restoration process and dual-modality representation learning. To ensure that the image and sequence branches generate consistent results, we propose the Image Feature Recombination module (IFR) and differentiable rasterization loss. Moreover, we also design a pre-training strategy using uni-modality data to relieve the lack of dual-modality data.

In summary, the contributions of this paper are fourfold:

- We analyze the conflict between the consistency constraints implicit in the prior knowledge of Chinese characters and the aesthetic designing of calligraphy fonts. To resolve this conflict, we propose a novel few-shot Chinese font generation model, DeepCalliFont, effectively integrating the information of glyph images and writing trajectories via a dual-modality representation learning.
- We propose an Image Feature Recombination module and a differentiable rasterization loss to ensure the consistency of two modalities, further improving the quality of synthesized glyph images.
- We design a pre-training strategy based on the coordinated representation of our DeepCalliFont, compensating for the lack of dual-modality data by exhaustively exploiting available uni-modality data.
- Quantitative and qualitative experiments demonstrate the superiority of our DeepCalliFont in the tasks of both regular and calligraphy font synthesis to existing methods, and the effectiveness of the proposed modules.

Related Work

Glyph Image Synthesis

Glyph image synthesis is a sub-task of image synthesis. Early attempts applied image generation methods to synthesize glyph images. For example, Tian (Tian 2017) proposed *zi2zi* by modifying *pix2pix* (Isola et al. 2017). Jiang et al. (2017) improved *zi2zi* and designed DCFont, an end-to-end font style transfer system, to generate Chinese font with less online data. For the few-show font generation task, Azadi et al. (2018) proposed an end-to-end stacked conditional GAN model, MC-GAN. Zhang et al. (2018) used two encoders to extract content and style features, and applied skip-connections to ensure glyph correctness. Gao et al. (2019) designed AGIS-Net with three discriminators to transfer both glyph shape and texture. There also exist some methods that utilize a large number of unpaired data to avoid time-consuming annotation. Wen et al. (2021) designed an end-to-end Chinese calligraphy font generation framework, ZiGAN, utilizing unpaired data to catch fine-grained features. Xie et al. (2021) proposed an unsupervised method, DGFont, whose key idea is using deformable convolutions to properly transfer the source glyph features to the features of target glyph images.

To synthesize better glyph details, some researchers also sought the help of high-level prior knowledge of Chinese characters. For example, SA-VAE (Sun et al. 2017) and CalliGAN (Wu, Yang, and Hsu 2020) added the stroke or component encoding vectors into the feature. Zeng et al. (2021), Huang et al. (Huang et al. 2020) and Kong et al. (2022) designed stroke-aware or component-aware discriminators. Based on the same components in characters, Park et al. proposed LF-Font (Park et al. 2020) and MX-Font (Park et al. 2021), establishing a relation between the local features and components. Similarly, Tang et al. (2022) developed an algorithm to search more suitable reference sets based on the prior knowledge of the existence of repeating components in FS-Font. Although these existing methods can synthesize visually-pleasing glyph images, they typically fail to synthesize satisfactory calligraphy fonts, which do not meet the consistency assumptions in prior knowledge.

Writing Trajectory Synthesis

The writing trajectory synthesis task aims to generate the writing trajectories of glyphs in target font styles. Compared with glyph images, writing trajectory synthesis is more challenging. Because there is no brush rendering, the trajectory style is mainly reflected by global structures and connected strokes. Generally, writing trajectories represent as a sequence of key points with drawing control labels. Therefore, some researchers used sequential models (e.g., RNN and LSTM) to generate target font trajectories. For example, Ha (2015) used RNN to synthesize writing trajectories. Zhang et al. (2017) also used RNN to identify and generate Chinese character trajectories and proposed a pre-processing method to make data suitable to sequential models. Tang et al. (2019) proposed an RNN-based network with a monotonic attention mechanism. And later, they proposed Write-LikeYou (Tang and Lian 2021) with two BiLSTM encoders and attention modules.

On the other hand, there also exist some approaches that deform the source trajectory sequence to generate the target trajectory sequence. For example, Miyazaki et al. (2019) proposed a stroke extraction method and calculated the deformation matrix from the source skeleton to the target one. Liu and Lian (Liu and Lian 2021) proposed FontRL, using reinforcement learning to perform multiple TPS transformations on the stroke skeleton and synthesizing the glyph skeleton. FontRL performs satisfactorily but requires that the source skeleton has the same number of strokes as the target. Thereby, it is unsuited for synthesizing calligraphy fonts with omitted strokes.

A similar task is vector font synthesis (Lopes et al. 2019; Carlier et al. 2020; Smirnov et al. 2020; Reddy et al. 2021; Wang and Lian 2021), which synthesizes drawing commands instead of writing trajectories. These methods can output vector font libraries directly. However, it is still hard for existing methods to directly synthesis high-quality vector glyphs for complex Chinese characters (Aoki and Aizawa 2022; Wang et al. 2023). Therefore, we choose writing trajectories as the sequence modality in our work.

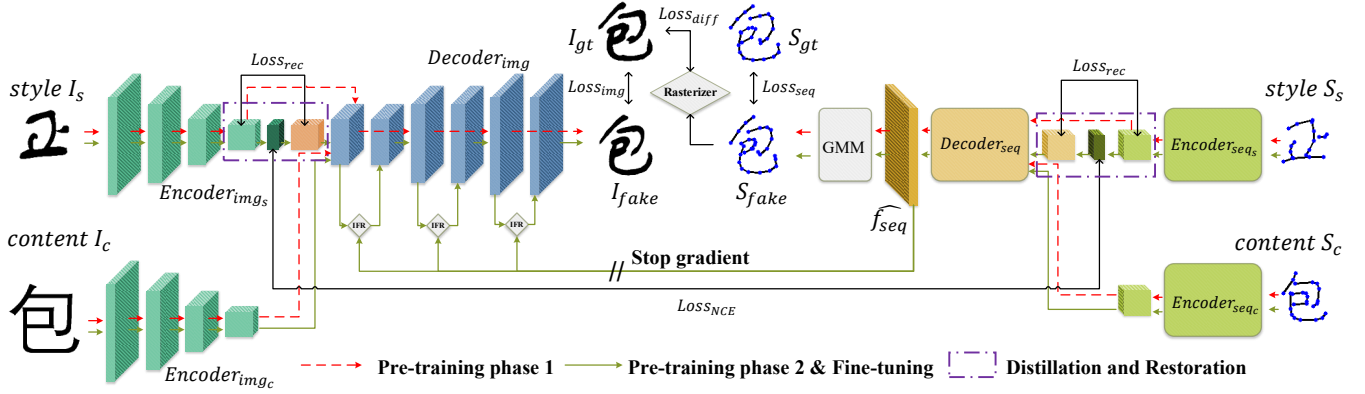


Figure 2: An overview of our model consisting of an image branch and a sequence branch. We first train two branches separately with uni-modality data (red dot lines) and then train the entire network jointly with dual-modality data (green solid lines).

Multi-Modality Font Generation

A Glyph can be represented in many modalities, such as the rasterization image, writing trajectory, glyph layout, components, etc. Some methods integrate two or more modalities to improve the synthesized font quality. For example, Wang et al. (2021) proposed DeepVecFont concatenating features of rasterization images and vector glyphs. SE-GAN (Yuan et al. 2022) used skeleton images to guide glyph image synthesis. Liu et al. (2022) designed XMP-Font integrating rasterization images and stroke order with cross-attention modules. Unlike the above methods using the “joint representation” (Baltrušaitis, Ahuja, and Morency 2018), we choose the “coordinated representation” (Baltrušaitis, Ahuja, and Morency 2018) because the glyph image contains information about other modalities. To be specific, we map dual-modality features to the same domain via a distillation-restoration process instead of concatenating them.

Method

As shown in Fig. 2, DeepCalliFont is an end-to-end trainable network, consisting of two symmetrical sub-networks, the image branch and the sequence branch. In this section, we first introduce the data representation in our model. Then, we describe the image and sequence branches in detail, respectively, and present the dual-modality representation learning used to map features of two modalities into the same domain. Finally, we briefly introduce the pre-training strategy of our method.

Data Representation

The proposed DeepCalliFont processes and synthesizes both glyph images (128×128) and the corresponding writing trajectories (see Fig. 3). We denote the content, style, and target glyph images as I_c , I_s , and I_{gt} , respectively. For the sequence modality, we denote the writing trajectory of a character as a sequence of key points $[P_1, P_2, P_3, \dots, P_n]$. Some existing methods use format-3 (Graves 2013) or format-5 (Ha 2015; Tang et al. 2019; Tang and Lian 2021) to represent the key point P_i , which consists of a point coordinate and a control label. As shown in Fig. 3, to distinguish




Figure 3: The two modalities used in DeepCalliFont: glyph images and writing trajectories. The red circle highlights a connected stroke in glyph image, so the format-6 representation of the corresponding point is $(x, y, 0, 0, 1, 0)$ (i.e., the control label $p_3 = 1$).

whether there is a connected stroke between strokes or not, we extend the original format-5 (Tang and Lian 2021) into a new data format, format-6, which can be represented as $(x, y, p_1, p_2, p_3, p_4)$, where $x, y \in [-1, 1]$ denote the point coordinate values, and (p_1, p_2, p_3, p_4) is a one-hot control label. $p_1 = 1$ means that there should be a visible line segment between this point and the next one; $p_2 = 1$ denotes the end of a stroke with no connection to the next stroke; $p_3 = 1$ means the end of a stroke that has a visual connection with the next stroke; and $p_4 = 1$ means the end of the writing process.

Sequence Branch

The sequence branch is a transformer-based seq2seq model, which can be formulated as:

$$\begin{aligned}
 f_{seq_s} &= \frac{1}{L} \sum_{i=1}^L Encoder_{seq_s}(S_s)^i; \\
 f_{seq_c} &= Encoder_{seq_c}(S_c); \\
 f_{seq} &= Decoder_{seq}(f_{seq_c}, f_{seq_s}); \\
 S_{fake-point} &= GMM(f_{seq}); \\
 S_{fake-label} &= ControlLabelClassifier(f_{seq});
 \end{aligned} \tag{1}$$

where S_s and S_c denote the style and content reference sequences, respectively, and L is the sequence length.

Following existing works, we also use Gaussian Mixture Model (GMM) (Tang et al. 2019; Tang and Lian 2021) to model the coordinates of key points as M bivariate normal distributions, and predict the control label by a control label classifier, which is a fully connected layer. The points and control labels can be represented as:

$$S_{fake-point} = \{\pi^i, \mu_x^i, \mu_y^i, \sigma_x^i, \sigma_y^i, \rho_{xy}^i\}_{i=1}^M; \quad (2)$$

$$S_{fake-label} = [q_1, q_2, q_3, q_4]; \quad (3)$$

where π denotes the probability of M distributions, q denotes the predicted control labels, and $M = 20$ by default.

The losses of the GMM and control label classifier are defined as:

$$loss_{point} = -\frac{1}{L} \sum_{i=1}^L \log(p(x, y)); \quad (4)$$

$$p(x, y) = \sum_{i=1}^M \pi^i \mathcal{N}(x, y | \mu_x^i, \mu_y^i, \sigma_x^i, \sigma_y^i, \rho_{xy}^i); \quad (5)$$

$$loss_{label} = CE(q, label_{control}); \quad (6)$$

where CE means the cross entropy loss, and $label_{control}$ denotes the target control label.

Differentiable Rasterizer Loss. To generate the writing trajectory located in the target glyph image, we proposed a differentiable rasterization loss based on Unsigned Distance Field (UDF). UDF defines a distance field between a grid and an object. The object in the sequence branch is a writing trajectory, which is represented as the set of all line segments between the key points. So the UDF in DeepCalliFont can be formalized as follows:

$$Dist(x, L) = \min_{l \in L} (Dist(x, l)), \quad x \in grid; \quad (7)$$

$$grid = \{(i, j) | i \in \{0, \dots, H-1\}, j \in \{0, \dots, W-1\}\}; \quad (8)$$

$$Dist(x, l) = \begin{cases} |x\vec{p}_s|, & p_s\vec{p}_t \cdot p_s\vec{x} < 0, \\ |x\vec{p}_t|, & p_t\vec{p}_s \cdot p_t\vec{x} < 0, \\ \frac{abs(x\vec{p}_s \times x\vec{p}_t)}{|p_s\vec{p}_t|}, & otherwise; \end{cases} \quad (9)$$

where L denotes the writing trajectory, p_s and p_t are the starting and ending points of the line segment l . Moreover, the distance $Dist(x, l)$ is set to ∞ if the line l is invisible (i.e., $p_1 = 0$ in format-6).

After getting the UDF of the synthesized writing trajectory, we render it into a raster image. To ensure the differentiability of the rendering process, we use the *Sigmoid* function to quantify the pixel value and get the image I_{diff} . The rendering process can be formalized as follows:

$$I_{diff} = \{1 - Sigmoid(\theta(Dist(x, S_{fake-point}) - w)) | x \in grid\}, \quad (10)$$

where w denotes the line width, and θ is a hyper-parameter. In this paper, θ and w are chosen as 100 and 2, respectively.

Then, we design a loose differentiable rasterizer loss, which is defined as:

$$loss_{diff} = ||ReLU(I_{diff} - I_{gt})||_2^2. \quad (11)$$

Loss Function. Finally, we calculate the complete loss function of the sequence branch by:

$$loss_{seq} = \lambda_1 loss_{point} + \lambda_2 loss_{label} + \lambda_3 loss_{diff}, \quad (12)$$

where λ_1 , λ_2 , and λ_3 are hyper-parameters that control the weights of $loss_{point}$, $loss_{label}$ and $loss_{diff}$.

Image Branch

The image branch is a CNN-based auto-encoder, including a style encoder, a content encoder, and a glyph decoder with skip connections between the content encoder and the decoder. The image branch can be formulated as:

$$\begin{aligned} f_{img_s} &= Encoder_{img_s}(I_s); \\ f_{img_c} &= Encoder_{img_c}(I_c); \\ I_{fake} &= Decoder_{img}(f_{img_s}, f_{img_c}); \end{aligned} \quad (13)$$

where I_s and I_c denote the style and content reference images, respectively.

Image Feature Recombination Module. Since the image feature aims to the global glyph while the sequence feature captures information of local writing details, it is possible to integrate these two complementary features to create a more discriminative one. Therefore, we design the Image Feature Recombination (IFR) module to recombine the image feature under the guidance of the sequence feature through an attention module. Because image features and sequence features are in different distributions (i.e., $f_{img} \sim p(img)$ and $f_{seq} \sim p(seq)$), we utilize two attention modules to recombine the image feature instead of directly concatenating these two types of features in the IFR module. In this manner, we can keep the recombined image feature in the same domain as the original image feature.

As shown in Fig. 4, we first map the image and sequence features to query (Q), key (K), and value (V) vectors through linear layers:

$$\begin{aligned} Q_{img} &= Linear_{Q_{img}}(f_{img}^i)^T \in R^{(H \times W) \times d}, \\ K_{img} &= Linear_{K_{img}}(f_{img}^i)^T \in R^{(H \times W) \times d}, \\ K_{seq} &= Linear_{K_{seq}}(f_{seq}) \in R^{L \times d}, \\ V_{seq} &= Linear_{V_{seq}}(f_{seq}) \in R^{L \times d}, \end{aligned} \quad (14)$$

where f_{img}^i denotes the feature outputted from the i th layer of the image decoder, f_{seq} denotes the feature generated by the sequence decoder, $H \times W$ is the size of image feature, L is the length of sequence feature, and d is the feature size.

Then, we calculate Q_{seq} through an attention module and layer normalization:

$$Q_{seq} = LayerNorm(\text{softmax}(\frac{Q_{img}K_{seq}^T}{\sqrt{d}})V_{seq}). \quad (15)$$

After getting the feature Q_{seq} , we recombine the original image feature f_{img}^i by another attention module:

$$f_{img}^i{}' = \text{softmax}(\frac{Q_{seq}K_{img}^T}{\sqrt{d}})f_{img}^i. \quad (16)$$




Figure 4: The illustration of the proposed Image Feature Recombination module. To keep the original feature f_{img}^i and the recombined feature f_{img}^i in the same domain, we design an attention module mapping the feature from image to sequence domain and another attention module mapping the feature from sequence to image domain.

Under the guidance of the sequence feature, we get a new feature f_{img}^i by recombining image features f_{img}^i extracted from different layers. We assume that the pixels in an image are i.i.d. and the convolution layer uses a fixed kernel to extract the feature from the image, so the features are also i.i.d. in the spatial level. Moreover, the attention mechanism is a linear combination of the value vector. Therefore, the IFR module keeps the recombined image feature f_{img}^i in the same domain as the original image feature f_{img}^i .

In addition, we stop the gradient of sequence feature f_{seq} to avoid that the unconverged image branch in the early training phase negatively influences the sequence branch.

Loss Function. We use MAELoss and the perceptual loss (2016) to optimize the image branch:

$$\begin{aligned} loss_{pixel} &= |I_{gt} - I_{fake}|; \\ loss_{percep} &= |Gram(vgg(I_{gt})) - Gram(vgg(I_{fake}))|; \\ loss_{img} &= \lambda_4 loss_{pixel} + \lambda_5 loss_{percep}; \end{aligned} \quad (17)$$

where vgg denotes a pre-trained VGG network, $Gram$ means the Gram matrix of features, λ_4 and λ_5 are hyper-parameters that control the weights of two losses.

Dual-Modality Representation Learning

As above mentioned, image and sequence features correspond to different feature domains. Therefore, we map style features f_{img_s} and f_{seq_s} into the same domain through **distillation** (dimension reduction) and **restoration** (dimension increasement). Specifically, we first distillate the feature by a linear layer and normalize it. Then we restore the feature via another linear layer to reconstruct the original feature.

The distillation-restoration process can be formulated as:

$$\begin{aligned} f_{d_x} &= \frac{Linear_d(f_{x_s})}{|Linear_d(f_{x_s})|}, \quad x \in \{img, seq\}; \\ f_{r_x} &= Linear_r(f_{d_x}), \quad x \in \{img, seq\}; \end{aligned} \quad (18)$$

where the features f_{x_s} and f_{r_x} have the same dimension, and the dimension of f_{d_x} is the half of theirs.

After getting the distilled features f_{d_x} , we apply the InfoNCE loss (Oord, Li, and Vinyals 2018) to the distilled image and sequence features. Because the feature vectors are normalized, we can measure the similarity with the cosine distance. The InfoNCE loss is computed by:

$$\begin{aligned} sim(img_i, seq_j) &= \cos(f_{d_{img_i}}, f_{d_{seq_j}}); \\ loss_{NCE} &= \frac{1}{B} \sum_{i=1}^B -\log \frac{\exp(\tau \cdot sim(img_i, seq_i))}{\sum_{j=1}^B \exp(\tau \cdot sim(img_i, seq_j))}; \end{aligned} \quad (19)$$

where f_{d_x} denotes the distilled feature, τ denotes the learnable temperature, and B means the batch size.

After restoration, we feed f_{r_x} into the decoder instead of f_{x_s} . To keep the consistency between the reconstructed feature f_{r_x} and the original f_x , we also adopt a feature reconstruction loss:

$$loss_{rec} = \|f_{img_s} - f_{r_{img}}\|_2^2 + \|f_{seq_s} - f_{r_{seq}}\|_2^2. \quad (20)$$

Moreover, we apply deep metric learning (Aoki et al. 2021) to the features f_{x_s} and f_{d_x} , which can be formalized as:

$$\begin{aligned} loss_{dml} &= \\ &CE(D_1(f_{img_s}), label) + CE(D_2(f_{d_{img}}), label) + \\ &CE(D_3(f_{seq_s}), label) + CE(D_4(f_{d_{seq}}), label), \end{aligned} \quad (21)$$

where D denotes the style classifier (a fully connected layer), and $label$ denotes the ground-truth font style.

We map the image and sequence features into the same domain through distillation and restoration. In this manner, the recombined feature can contain dual-modality information. In summary, the complete loss function of our Deep-CalliFont is defined as:

$$\begin{aligned} loss &= loss_{img} + loss_{seq} + \\ &loss_{NCE} + loss_{rec} + loss_{dml}. \end{aligned} \quad (22)$$

Pre-training Strategy

Considering that it is hard to build large-scale dual-modality datasets, we design a two-phase pre-training strategy. As shown in Fig. 2, in phase 1, we train the image and sequence branches separately using a large amount of uni-modality data. In phase 2, we add modality interaction modules, including the IFR module, the differentiable rasterizer loss, and the distillation-restoration module, and jointly train two branches using a small amount of dual-modality data.

Experiments

Experimental Settings

We use 251 fonts and CASIA Online Chinese Handwriting Databases (Liu et al. 2011) to pre-train two branches

Content	基蝠巡扒扣垢牲百衍塌
MX-Font	基蝠巡扒扣垢牲百衍塌
AGIS-Net	基蝠巡扒扣垢牲百衍塌
WLY (Seq)	基蝠巡扒扣垢牲百衍塌
Ours	Seq 基蝠巡扒扣垢牲百衍塌
	Image 基蝠巡扒扣垢牲百衍塌
GT	基蝠巡扒扣垢牲百衍塌

Figure 5: Comparison of calligraphy font synthesis results obtained by DeepCalliFont and other three methods. For the proposed DeepCalliFont, we show the synthesis results of two branches. We highlight the synthesis results of connected strokes with boxes. Moreover, the red lines in synthesized sequences denote the connected strokes.

separately in the pre-training phase 1. Then, in the pre-training phase 2, we selected 42 fonts used in (Jiang et al. 2019), each of which consists of 3,000 glyph images and their corresponding writing trajectories, to train the whole model. Since there is no connected stroke annotation available in the above-mentioned datasets, we automatically generate pseudo annotations as follows. We define two strokes as connected if and only if the distance between the previous stroke’s endpoint and the next stroke’s start point is less than a given threshold and the distance between corresponding points in the mean skeleton is greater than the threshold.

To demonstrate the generalizability of the proposed DeepCalliFont, we construct two test datasets, i.e., a regular font test set (without connected strokes) and a calligraphy font test set (with connected strokes). We use 30 fonts in the dataset collected by Liu and Lian (Liu and Lian 2023) as the regular font test set. We also collect ten calligraphy fonts as the other test set. In our experiments, we fine-tune networks on 100 sample characters and test them on other 6,663 Chinese characters. Then, we evaluate the quality of synthesized glyph images based on MAE, FID (Heusel et al. 2017), and LPIPS (Zhang et al. 2018), and calculate the similarity between generated and ground-truth sequences via dynamic time warping (DTW) (Bahdanau, Cho, and Bengio 2014).

Quantitative and Qualitative Comparison

We compare the performance of our DeepCalliFont with seven existing methods, including zi2zi (Tian 2017), EMD (Zhang, Zhang, and Cai 2018), AGIS-Net (Gao et al. 2019), MX-Font (Park et al. 2021), DG-Font (Xie et al. 2021), FontRL (Liu and Lian 2021), and WriteLikeYou (WLY) (Tang and Lian 2021). Since zi2zi, EMD, DGFont perform poorly, and FontRL requires the same number of strokes, we compare our method with the others on calligraphy fonts. For the sake of fair comparison, we pre-train and fine-tune all networks with the same data, including uni-modality and dual-modality data.

Tab. 2 shows that our DeepCalliFont outperforms other

Content	茁搔封臚斥沪蚱芥届捏
Zi2zi	茁搔封臚斥沪蚱芥届捏
EMD	茁搔封臚斥沪蚱芥届捏
AGIS-Net	茁搔封臚斥沪蚱芥届捏
MX-Font	茁搔封臚斥沪蚱芥届捏
DG-Font	茁搔封臚斥沪蚱芥届捏
FontRL	茁搔封臚斥沪蚱芥届捏
WLY (Seq)	茁搔封臚斥沪蚱芥届捏
Ours	Seq 茁搔封臚斥沪蚱芥届捏
	Image 茁搔封臚斥沪蚱芥届捏
GT	茁搔封臚斥沪蚱芥届捏

Figure 6: Comparison of regular font synthesis results generated by DeepCalliFont and other methods.

	gt	Uni-modality	Dual-modality		gt	Uni-modality	Dual-modality	
	Seq	image	Seq	image	Seq	image	Seq	image
	捌捌捌捌捌				熬熬熬熬熬			
	次次次次次				哥哥哥哥哥			

Figure 7: Comparison of synthesis results obtained by the uni-modality branch or the dual-modality model. The results of our dual-modality model have better glyph consistency and better details.

methods in all metrics. The proposed approach exhibits a significant advantage against them, especially in the metrics of FID and LPIPS. As shown in Fig. 5, MX-Font and AGIS-Net struggle to synthesize glyphs with connected strokes, although they use adversarial learning or additional prior knowledge. WriteLikeYou solves this problem to some extent, but it often fails to satisfactorily reconstruct the whole glyph structure. On the contrary, our DeepCalliFont captures the difference on strokes between the source font and the target font, and synthesizes the highest-quality glyph images in calligraphy styles. Since our model imposes several constraints to keep the features of two modalities consistent, it can synthesize consistent glyph images and writing trajectories that coincide well.

Although DeepCalliFont is specifically designed to handle calligraphy fonts consisting of glyphs that often contain connections between strokes, its performance on regular fonts is also satisfactory. Synthesis results of different methods are shown in Fig. 6, from which we can see that our DeepCalliFont still outperforms other state-of-the-art methods in the task of regular font synthesis.

Image	Seq	IFR	Distillation	Diff	NCE Loss	DML	Uni-modality pre-train	MAE↓	FID↓	LPIPS↓
✓	✓	✓	✓	✓	✓	✓	✓	0.1254	14.6	0.1717
✓	✓		✓	✓	✓	✓	✓	0.1270	23.7	0.1882
✓	✓	✓		✓	✓	✓	✓	0.1311	24.6	0.1969
✓	✓		✓		✓	✓	✓	0.1285	24.6	0.1925
✓	✓	✓		✓		✓	✓	0.1279	27.0	0.1961
✓	✓	✓	✓	✓	✓		✓	0.1300	25.4	0.1943
✓						✓	✓	0.1348	21.4	0.1821
✓	✓	✓	✓	✓	✓	✓		0.1355	43.4	0.2197

Table 1: Ablation study results. Image: image branch; Seq: sequence branch; IFR: Image Feature Recombination module; Distillation: “Distillation-restoration” module; Diff: differentiable rasterization loss; NCE Loss: the InfoNCE loss (Eq. 19); DML: Deep metric learning (Eq. 21); Uni-modality pre-train: the pre-training stage 1.

Calligraphy Fonts				
Method	MAE↓	FID↓	LPIPS↓	
AGIS-Net	0.1383	463.9	0.2857	
MX-Font	0.1850	235.1	0.2754	
WriteLikeYou	0.1757	167.2	0.2714	
ours	0.1254	14.6	0.1718	
Regular Fonts				
Method	MAE↓	FID↓	LPIPS↓	DTW↓
zi2zi	0.1719	137.1	0.4050	-
EMD	0.1538	184.5	0.4628	-
AGIS-Net	0.1354	157.7	0.4584	-
MX-Font	0.2115	137.4	0.5055	-
DG-Font	0.1886	150.2	0.4878	-
FontRL	0.1350	40.7	0.2588	15.304
WriteLikeYou	0.2223	114.4	0.4565	5.920
ours	0.1242	105.9	0.1837	1.347

Table 2: Quantitative results of our DeepCalliFont and other methods on calligraphy and regular fonts.

Ablation Studies

In this subsection, we analyze the role of each proposed module or loss function. Tab. 1 shows the quantitative results demonstrating that those proposed modules and loss functions all help to improve the performance of our method.

Uni-Modality vs Dual-Modality Synthesis. Tab. 1 shows that our full model performs better in all metrics than the model using only the image branch. As we can see from Fig. 7, with the help of dual-modality representation learning and other proposed modules, our model is capable of synthesizing better glyph details.

Differentiable Rasterization Loss. The differentiable rasterization loss $loss_{diff}$ constrains the writing trajectories to be located inside the corresponding glyphs. From Fig. 8 where the synthesized trajectories are drawn on the target glyph images, we can see that the writing trajectory synthesized with $loss_{diff}$ matches better with the target image.

Image Feature Recombination Module. The Image Feature Recombination (IFR) module aims to improve the result of the image branch with the help of sequence features. Empirically, the sequence branch converges faster and mostly generates correct glyph structures. We can see from



Figure 8: Comparison of synthesis results of our method with or without the differentiable rasterizer loss, which limits the generated trajectory inside the corresponding glyph.




Figure 9: Comparison of synthesis results obtained by our method with or without the proposed IFR module. Through the IFR module, the image branch improves the details of strokes under the guidance of sequence features.

Fig. 9 that the image branch with IFR generates glyph images more consistent with the synthesized writing trajectories than the one without the IFR module. Moreover, the sequence feature can guide image branches to generate better details, which are difficult to correct through loss functions.

“Distillation-Restoration” Module. Due to the above-mentioned pre-training strategy, the two branches were designed to be independent. Therefore, we design the “distillation-restoration” module to combine these two branches into a unified network. To verify the effectiveness of this module, we remove it and apply the InfoNCE loss directly on features f_{img_s} and f_{seq_s} (see Tab. 1).

Deep Metric Learning. Following Aoki et al. (2021), we add deep metric learning into DeepCalliFont, helping our model distinguish font styles. Therefore, the $loss_{dml}$ mainly improves the method performance in of FID and LPIPS.

Two-Stage Pre-training Strategy In Tab. 1, we compare the results obtained with and without the first stage pre-training, validating its effectiveness.

Limitation

Since the few-shot font generation task is essentially ill-posed, synthesized images may differ from the target ground truth but are visually pleasing. More discussions and failure cases can be found in the supplementary material.

Conclusion

In this paper, we proposed a novel end-to-end few-shot Chinese font synthesis method, DeepCalliFont. To handle the connected strokes frequently contained in calligraphy fonts, we designed a dual-modality (i.e., glyph images and writing trajectories) few-shot font generation network. Specifically, we utilized contrastive learning to keep the consistency of image and sequence branches at the feature level. We also proposed the Image Feature Recombination module and differentiable rasterization loss to generate consistent and high-quality results. Meanwhile, we designed a pre-training strategy to use a large amount of uni-modality data to compensate for the lack of dual-modality data. Both quantitative and qualitative experiments showed that the proposed DeepCalliFont can synthesize higher-quality calligraphy fonts with more reasonable connected strokes than other SOTA methods and also perform satisfactorily on regular fonts.

Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No.: 62372015), Center For Chinese Font Design and Research, and Key Laboratory of Intelligent Press Media Technology.

References

- Aoki, H.; and Aizawa, K. 2022. SVG Vector Font Generation for Chinese Characters with Transformer. *arXiv preprint arXiv:2206.10329*.
- Aoki, H.; Tsubota, K.; Ikuta, H.; and Aizawa, K. 2021. Few-Shot Font Generation with Deep Metric Learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 8539–8546. IEEE.
- Azadi, S.; Fisher, M.; Kim, V. G.; Wang, Z.; Shechtman, E.; and Darrell, T. 2018. Multi-content gan for few-shot font style transfer. In *CVPR*, 7564–7573.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Baltrušaitis, T.; Ahuja, C.; and Morency, L.-P. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2): 423–443.
- Carlier, A.; Danelljan, M.; Alahi, A.; and Timofte, R. 2020. Deepsvg: A hierarchical generative network for vector graphics animation. *Advances in Neural Information Processing Systems*, 33: 16351–16361.
- Gao, Y.; Guo, Y.; Lian, Z.; Tang, Y.; and Xiao, J. 2019. Artistic glyph image synthesis via one-stage few-shot learning. *ACM Transactions on Graphics (TOG)*, 38(6): 1–12.
- Graves, A. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- Ha, D. 2015. Recurrent Net Dreams Up Fake Chinese Characters in Vector Format with TensorFlow. *blog.otoro.net*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 6626–6637.
- Huang, Y.; He, M.; Jin, L.; and Wang, Y. 2020. RD-GAN: few/zero-shot Chinese character style transfer via radical decomposition and rendering. In *European conference on computer vision*, 156–172. Springer.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *CVPR*, 1125–1134.
- Jiang, Y.; Lian, Z.; Tang, Y.; and Xiao, J. 2017. DCFont: an end-to-end deep Chinese font generation system. In *SIGGRAPH Asia*, 22. ACM.
- Jiang, Y.; Lian, Z.; Tang, Y.; and Xiao, J. 2019. SCFont: Structure-guided Chinese Font Generation via Deep Stacked Networks.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 694–711. Springer.
- Kong, Y.; Luo, C.; Ma, W.; Zhu, Q.; Zhu, S.; Yuan, N.; and Jin, L. 2022. 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*, 13482–13491.
- Lian, Z.; Zhao, B.; Chen, X.; and Xiao, J. 2018. EasyFont: a style learning-based system to easily build your large-scale handwriting fonts. *TOG*, 38(1): 1–18.
- Liu, C.-L.; Yin, F.; Wang, D.-H.; and Wang, Q.-F. 2011. CA-SIA online and offline Chinese handwriting databases. In *ICDAR*, 37–41. IEEE.
- Liu, W.; Liu, F.; Din, F.; He, Q.; and Yi, Z. 2022. XMP-Font: Self-Supervised Cross-Modality Pre-training for Few-Shot Font Generation. *arXiv preprint arXiv:2204.05084*.
- Liu, Y.; and Lian, Z. 2021. FontRL: Chinese Font Synthesis via Deep Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2198–2206.
- Liu, Y.; and Lian, Z. 2023. FontTransformer: Few-shot high-resolution Chinese glyph image synthesis via stacked transformers. *Pattern Recognition*, 141: 109593.
- Lopes, R. G.; Ha, D.; Eck, D.; and Shlens, J. 2019. A learned representation for scalable vector graphics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7930–7939.
- Miyazaki, T.; Tsuchiya, T.; Sugaya, Y.; Omachi, S.; Iwamura, M.; Uchida, S.; and Kise, K. 2019. Automatic generation of typographic font from small font subset. *IEEE computer graphics and applications*, 40(1): 99–111.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.

- Park, S.; Chun, S.; Cha, J.; Lee, B.; and Shim, H. 2020. Few-shot Font Generation with Localized Style Representations and Factorization. *arXiv preprint arXiv:2009.11042*.
- Park, S.; Chun, S.; Cha, J.; Lee, B.; and Shim, H. 2021. Multiple Heads are Better than One: Few-shot Font Generation with Multiple Localized Experts. *arXiv preprint arXiv:2104.00887*.
- Reddy, P.; Gharbi, M.; Lukac, M.; and Mitra, N. J. 2021. Im2vec: Synthesizing vector graphics without vector supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7342–7351.
- Smirnov, D.; Fisher, M.; Kim, V. G.; Zhang, R.; and Solomon, J. 2020. Deep parametric shape predictions using distance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 561–570.
- Sun, D.; Ren, T.; Li, C.; Su, H.; and Zhu, J. 2017. Learning to write stylized chinese characters by reading a handful of examples. *arXiv preprint arXiv:1712.06424*.
- Tang, L.; Cai, Y.; Liu, J.; Hong, Z.; Gong, M.; Fan, M.; Han, J.; Liu, J.; Ding, E.; and Wang, J. 2022. Few-Shot Font Generation by Learning Fine-Grained Local Styles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7895–7904.
- Tang, S.; and Lian, Z. 2021. Write Like You: Synthesizing Your Cursive Online Chinese Handwriting via Metric-based Meta Learning. In *Computer Graphics Forum*, volume 40, 141–151. Wiley Online Library.
- Tang, S.; Xia, Z.; Lian, Z.; Tang, Y.; and Xiao, J. 2019. FontRNN: Generating Large-scale Chinese Fonts via Recurrent Neural Network. In *Computer Graphics Forum*, volume 38, 567–577. Wiley Online Library.
- Tian, Y. 2017. zi2zi: Master chinese calligraphy with conditional adversarial networks, 2017. Retrieved Jun, 3: 2017.
- Wang, Y.; and Lian, Z. 2021. DeepVecFont: synthesizing high-quality vector fonts via dual-modality learning. *ACM Transactions on Graphics (TOG)*, 40(6): 1–15.
- Wang, Y.; Wang, Y.; Yu, L.; Zhu, Y.; and Lian, Z. 2023. DeepVecFont-v2: Exploiting Transformers to Synthesize Vector Fonts with Higher Quality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18320–18328.
- Wen, Q.; Li, S.; Han, B.; and Yuan, Y. 2021. ZiGAN: Fine-grained Chinese Calligraphy Font Generation via a Few-shot Style Transfer Approach. In *Proceedings of the 29th ACM International Conference on Multimedia*, 621–629.
- Wu, S.-J.; Yang, C.-Y.; and Hsu, J. Y.-j. 2020. Calligan: Style and structure-aware chinese calligraphy character generator. *arXiv preprint arXiv:2005.12500*.
- Xie, Y.; Chen, X.; Sun, L.; and Lu, Y. 2021. DG-Font: Deformable Generative Networks for Unsupervised Font Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5130–5140.
- Yuan, S.; Liu, R.; Chen, M.; Chen, B.; Qiu, Z.; and He, X. 2022. SE-GAN: Skeleton Enhanced GAN-based Model for Brush Handwriting Font Generation. *arXiv preprint arXiv:2204.10484*.
- Zeng, J.; Chen, Q.; Liu, Y.; Wang, M.; and Yao, Y. 2021. Strokegan: Reducing mode collapse in chinese font generation via stroke encoding. In *proceedings of AAAI*, volume 3.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.
- Zhang, X.-Y.; Yin, F.; Zhang, Y.-M.; Liu, C.-L.; and Bengio, Y. 2017. Drawing and recognizing chinese characters with recurrent neural network. *IEEE transactions on pattern analysis and machine intelligence*, 40(4): 849–862.
- Zhang, Y.; Zhang, Y.; and Cai, W. 2018. Separating style and content for generalized style transfer. In *CVPR*, 8447–8455.