

# HF-FONT: FEW-SHOT FONT GENERATION VIA HIGH-FREQUENCY STYLE ENHANCEMENT AND FUSION

005 **Anonymous authors**

006 Paper under double-blind review

## ABSTRACT

011 Few-shot font generation aims to create new fonts with a limited number of glyph  
 012 references. It can be used to greatly reduce the workload of manual font design.  
 013 However, although existing methods have achieved satisfactory performance, they  
 014 still struggle to capture delicate glyph details, thus resulting in stroke errors, arti-  
 015 facts, and blurriness. To address these problems, we propose HF-Font, a novel  
 016 framework that generates fonts with higher structural fidelity. Specifically, in-  
 017 spired by the observation that high-frequency information of character images  
 018 often contains distinct style patterns (*e.g.*, glyph topology and stroke variation),  
 019 we develop a novel style-enhanced module to improve the style extraction by  
 020 incorporating high-frequency features from reference images using a high-pass  
 021 filter. Then, for guiding the generation process, we design a Style-Content Fu-  
 022 sion Module (SCFM), which integrates the style features with a component-wise  
 023 codebook that encodes content semantics. Moreover, we also introduce a style  
 024 contrastive loss to better transfer high-frequency features. Extensive experiments  
 025 show that our HF-Font outperforms the state-of-the-art methods in both qualita-  
 026 tive and quantitative evaluations, demonstrating its effectiveness across diverse  
 027 font styles and characters. Our source code will be released soon.

## 1 INTRODUCTION

031 The task of few-shot font generation allows to transfer the font style from a source domain to a  
 032 target domain based on only a few reference images. It can greatly alleviate the burden of time-  
 033 consuming and labor-intensive manual design, especially for some character-rich languages like  
 034 Chinese, Japanese, or Korean. Therefore, few-shot font generation techniques can benefit many  
 035 critical applications, including logo design, ancient character restoration, and so on.

036 At present, with the rapid development of deep learning architectures, such as Convolutional Neu-  
 037 ral Networks (CNNs) (Chen et al., 2022), Generative Adversarial Networks (GANs) (Goodfellow  
 038 et al., 2014), and diffusion models (Ho et al., 2020), researchers have made great advances in cre-  
 039 ating gratifying fonts. A widely-adopted strategy for this task is the style-content disentanglement  
 040 paradigm, which decouples the content and style representations from the given content and refer-  
 041 ence images. These two representations are then combined and decoded to generate the target glyph.  
 042 Early approaches (Zhang et al., 2018b; Gao et al., 2020; Xie et al., 2021; Wang et al., 2023) mainly  
 043 model font style as a set of universal statistic features, and utilize the extracted representations to  
 044 encode stylistic information. In addition, witnessing the structure variations and local correlations  
 045 in font styles, some works (Park et al., 2021b; Tang et al., 2022; Kong et al., 2022; Chen et al., 2024)  
 046 also employ structure-aware representations, which typically decompose characters into different  
 047 components and acquire multiple style representations to boost the performance. Nevertheless, de-  
 048 spite the remarkable progress, existing methods still have several drawbacks. The highly diverse  
 049 and intricate nature of font styles often leads to obvious defects in synthesized results, like incom-  
 050 plete or unwanted strokes, anomalous blurriness, and artifacts. Besides, a core challenge of few-shot  
 051 font generation lies in how to accurately capture the target style from just a few reference images.  
 052 Previous works tend to exhibit a limited performance in reproducing stylistic details, due to their  
 053 insufficient style extraction ability.

To tackle the above challenges, our key idea revolves around leveraging high-frequency features to  
 enhance the style extraction. As shown in Figure 1, [these high-frequency components are defined by](#)

054	Characters	奏 综 尊 滋	轴 衷 安 呱	逐 著 窗 贻	挝 辨 豹 爺
055	High-frequency	奏 综 尊 滋	轴 衷 安 呱	逐 著 窗 贻	挝 辨 豹 爺
056	Characters	柏 棚 半 犀	堡 彪 别 彬	秉 荻 荦 犇	暴 悸 隅 宾
057	High-frequency	柏 棚 半 犀	堡 彪 别 彬	秉 荻 荢 犇	暴 悸 隅 宾
058	Characters	瘤 昂 祜 跋	窟 必 把 蕤	奥 班 柄 爰	北 版 鼻 蝙
059	High-frequency	瘤 昂 祜 跋	窟 必 把 蕤	奥 班 柄 爰	北 版 鼻 蝙
060	Characters				
061	High-frequency				
062					

Figure 1: **Some character samples and their corresponding high-frequency components.** We find that high-frequency components provide clearer contours and highlight distinctive style patterns. Zoom in to see details.

rapid spatial intensity transitions along stroke boundaries, rather than by differences in color. They highlight the overall contours of the glyphs and reveal essential style attributes such as topology, stroke variation, and line spacing. Hence, the incorporation of high-frequency information facilitates a more effective and robust style extraction process.

In light of this, we put forward a new end-to-end few-shot font generation method called HF-Font. To be specific, we first develop a style-enhanced module to process both style reference images and their corresponding high-frequency counterparts in parallel. Given that the reference images usually contain background noise, we introduce a gate mechanism to filter out the irrelevant information. Subsequently, the style features extracted from both branches are adaptively fused with the component-wise content representation through an innovative Style-Content Fusion Module (SCFM), which includes two Grouped Residual Attention Blocks (GRABs) and one output adapter. Apart from these, we also design a style contrastive loss to further obtain discriminative features, promoting both realism and diversity in the generated results. Consequently, our method yields superior and impressive stylization performance, highlighting its potential for practical applications.

To summarize, our major contributions are as follows:

- We propose HF-Font, a novel few-shot font generation framework. It incorporates a style-enhanced module to leverage high-frequency information from reference images, thereby strengthening structural integrity.
- We devise an inventive Style-Content Fusion Module (SCFM), which efficiently integrates style information with the component-level codebook. Furthermore, a style contrastive loss is also employed to facilitate the transfer of high-frequency features.
- Extensive experimental results verify that our HF-Font surpasses the leading models in both qualitative and quantitative evaluations, which demonstrates its effectiveness and generalization capability across diverse font styles and characters.

## 2 RELATED WORK

### 2.1 IMAGE-TO-IMAGE TRANSLATION

Image-to-image (I2I) translation is to convert an image from one domain to another while preserving its semantic content. Previous methods, such as Pix2pix (Isola et al., 2017) and CycleGAN (Zhu et al., 2017), mainly leverage GANs (Goodfellow et al., 2014) for supervised or unsupervised settings. Later, motivated by the human ability of inductive reasoning, FUNIT (Liu et al., 2019) employs Adaptive Instance Normalization (AdaIN) (Huang & Belongie, 2017) to fuse the encoded content and style features. Recently, diffusion models have emerged as a powerful generative technology. ILVR (Choi et al., 2021) achieves high-quality performance using only a pre-trained unconditional Denoising Diffusion Probabilistic Model (DDPM) (Ho et al., 2020). SCDM (Ko et al., 2024) applies stochastic perturbations to semantic maps and conditions I2I translation on the diffused labels. Notably, since font generation can be viewed as a special I2I translation task, many generic I2I translation approaches can be adaptively modified for font generation.

108  
109

## 2.2 FEW-SHOT FONT GENERATION

110  
111  
112  
113  
114  
115  
116  
117  
118  
119

Few-shot font generation intends to create a required font library with just a handful of glyph references. A classical tactic involves extracting content and style features from the input source and reference images, respectively. EMD (Zhang et al., 2018b) and AGIS-Net (Gao et al., 2020) disentangle the representations of style and content, and model each font as a universal descriptor. DG-Font (Xie et al., 2021) introduces a feature deformation skip connection to capture the glyph deformations. CF-Font (Wang et al., 2023) extends DG-Font through incorporating a content fusion module to narrow the gap between the source and target fonts. Later, NTF (Fu et al., 2023) formulates font generation as a continuous transformation via a neural transformation field (Mildenhall et al., 2020). FontDiffuser (Yang et al., 2024) implements a multi-scale content aggregation block along with a style contrastive refinement module to guide the whole framework.

120  
121  
122  
123  
124  
125  
126  
127  
128  
129  
130  
131  
132  
133  
134  
135  
136  
137

For the generation of highly-structured characters, some notable studies leverage prior domain knowledge, such as stroke decomposition and trajectory, to optimize the final results. SA-VAE (Sun et al., 2018) stands out as the first attempt to integrate radicals and spatial structures into the generative model. LF-Font (Park et al., 2021b) represents component-wise style feature via a low-rank matrix factorization. MX-Font (Park et al., 2021a) utilizes a multi-head design, with each head weakly supervised to extract distinct local concepts. **MX-Font++** (Wang et al., 2025) extends **MX-Font** by introducing **Heterogeneous Aggregation Experts (HAE)** to enhance feature extraction and decoupling. Afterwards, CG-GAN (Kong et al., 2022) utilizes a component predictor to assist the generator during adversarial training, while FsFont (Tang et al., 2022) establishes a character-reference mapping relationship and employs cross-attention to align the patch-level features. Diff-Font (He et al., 2024) infuses predefined embedding tokens into the condition diffusion model to support the sampling process. IF-Font (Chen et al., 2024) replaces the source image with the Ideographic Description Sequence (IDS) to control glyph semantics. Yet, these methods remain dependent on the labels of component categories. In contrast, VQ-Font (Pan et al., 2023) constructs a vector quantization-based encoder to automatically extract components. **DA-Font** (Chen et al., 2025) integrates a **Dual-Attention Hybrid Module** to optimize style transfer and relation-aware feature harmonization. Nevertheless, prevailing approaches still fall short in effectively encapsulating intricate structural features. This motivates us to harness high-frequency information and design a novel feature fusion strategy, as such information offers sharper style cues.

138

139  
140

## 3 METHODOLOGY

141  
142  
143

## 3.1 OVERALL SCHEME

144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155

Given a set of  $k$  reference images  $x = \{x_i\}_{i=1}^k$  and a content image  $I_c$ , our model aims to generate a character  $I_o$  that retains the same content with  $I_c$  and the same style with  $x$ . As illustrated in Figure 2 (a), the generator mainly includes eight parts. Among them, the content encoder  $E_c$  extracts the feature representation  $f_c$  from  $I_c$ . It is pre-trained via a Vector Quantized Variational Auto-Encoder (VQ-VAE) to acquire a component-wise codebook  $F_c$  as well. The reference images  $x$  are first processed by a high-pass filter to obtain their high-frequency images  $h = \{h_i\}_{i=1}^k$ . Next, the spatial and high-frequency style encoders,  $E_{ss}$  and  $E_{sh}$ , extract style features  $f_{spa} = \{f_{is}\}_{i=1}^k$  and  $f_{fre} = \{f_{ir}\}_{i=1}^k$  from  $x$  and  $h$ , respectively. After that,  $f_{spa}$  is further refined through a gate mechanism to yield  $\hat{f}_{spa} = \{\hat{f}_{is}\}_{i=1}^k$ , which is then fused with  $f_{fre}$  and  $F_c$  via our proposed SCFM to derive the style feature  $f_{sa}$ . Meanwhile, the Content Alignment Module (CAM) adaptively re-weights and aggregates  $\hat{f}_{spa}$  to get the aligned content feature  $f_{ca}$ . Finally, the decoder  $D_r$  outputs the result  $I_o$  by concatenating  $f_c$ ,  $f_{sa}$ , and  $f_{ca}$  as inputs.

156  
157  
158  
159  
160  
161

During the model’s training, a multi-task discriminator  $D_{s,c}$  is employed to distinguish between the real and generated images, as shown in Figure 2 (b). Moreover, to ensure that the generated glyph accurately reflects the reference style while preserving the input content structure, it also performs the classifications on each character, identifying both its style and content categories. The complete architectures of the generator and discriminator are provided in Appendix A.1. **Notably, the content encoder  $E_c$  and the codebook  $F_c$  remain frozen during the main training stage, while all other components are fully trainable.**

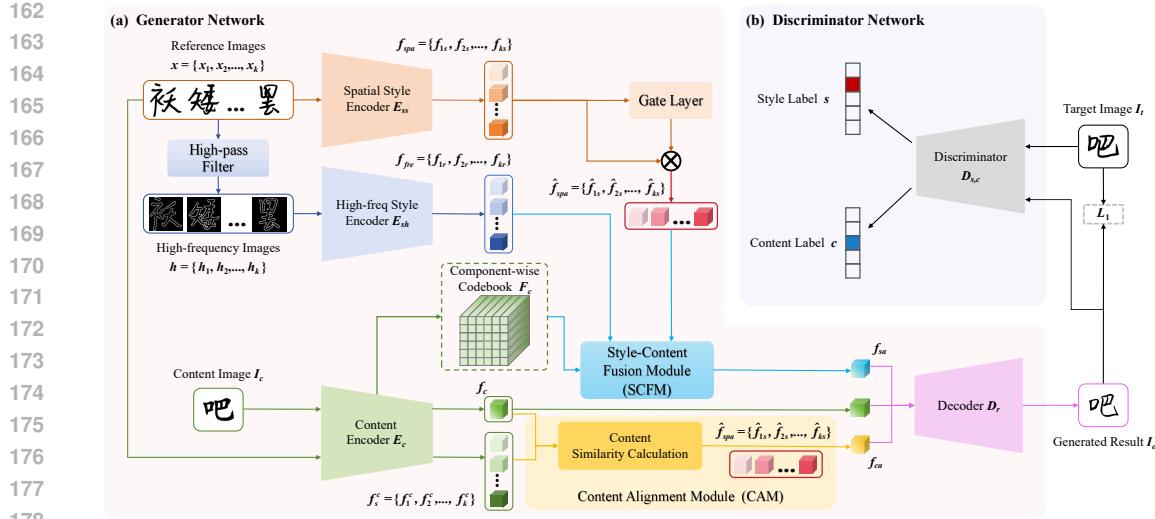


Figure 2: **Overview of our proposed HF-Font.** (a) The generator network mainly consists of the following parts: a pre-trained content encoder  $E_c$ , a spatial style encoder  $E_{ss}$ , a high-frequency style encoder  $E_{sh}$ , a high-pass filter, a gate mechanism, a Content Alignment Module (CAM), a Style-Content Fusion Module (SCFM), and a decoder  $D_r$ . (b) A discriminator network is used to distinguish the real and fake images, while also classifying the content and style categories of the generated characters.

### 3.2 STYLE-ENHANCED MODULE

To enrich the style representations, we introduce a style-enhanced module that exploits high-frequency information into the feature extraction process. In particular, a Laplacian kernel is used as a high-pass filter to obtain high-frequency counterparts  $h$  from reference images  $x$ . This operation efficiently emphasizes fine-grained details without relying on any frequency-domain transformation (Wahl, 2024). Next, two parallel style encoders,  $E_{ss}$  and  $E_{sh}$ , are employed to extract complementary style features  $f_{spa}$  and  $f_{fre}$  from  $x$  and  $h$  separately. Though structurally identical, these two encoders do not share weights with each other. Moreover, a gate mechanism is integrated into the spatial branch to modulate  $f_{spa}$ , allowing only informative style cues to pass.

**Gate Mechanism.** Intuitively, the stroke regions in reference images are often sparse, which makes style extraction susceptible to background noise (Dai et al., 2024). **Here, background noise refers to the non-informative or low-response activations in intermediate feature maps during convolutional encoding, rather than the white background of the glyph images.** To solve this problem, we bring in a gate mechanism to selectively filter the reference features. Specifically, the initial spatial style feature  $f_{spa} = \{f_{is}\}_{i=1}^k$  is fed into a gate layer, consisting of a learnable fully-connected layer followed by a Sigmoid activation, to produce the gate units  $w = \{w_i\}_{i=1}^k$ . Each  $w_i$  flexibly controls the information flow of  $f_{is}$ , where higher values indicate stronger retention of the style signals. The refined features  $\hat{f}_{spa} = \{\hat{f}_{is}\}_{i=1}^k$  are obtained by element-wise multiplication, i.e.,  $\hat{f}_{is} = w_i \cdot f_{is}$ . This design effectively suppresses extraneous noise while preserving meaningful style patterns.

### 3.3 GLYPH FEATURE DECOMPOSITION AND CONTENT ALIGNMENT MODULE

**Glyph Feature Decomposition.** In our HF-Font, the content encoder is pre-trained through a glyph feature decomposition network, which projects each character into a component-wise codebook. This network is trained on a certain character set for image reconstruction. As illustrated in Figure 3, the content encoder  $E_c$  is built on CNN and maps a character image  $I_f$  into latent representation  $Z_c$ . Then, vector quantization (Aaron et al., 2017) is applied to discrete  $Z_c$  as:

$$z_c^i = e_c^j, \quad \text{s.t.} \quad j = \arg \min_{j \in \{1, 2, \dots, d\}} \|z_c^i - e_c^j\|^2 \quad (1)$$

216 where each spatial vector  $z_c^i$  in  $Z_c$  is replaced  
 217 by its nearest code vector  $e_c^j$  from the codebook  
 218  $F_c$ , which contains  $d$  code vectors. Finally, the  
 219 reconstruction decoder  $D_r$  takes the matched  
 220 codes  $Z_q$  as input to reconstruct the glyph  $I_r$ .  
 221 Visual analyses of the codebook are presented  
 222 in Appendix A.2.

223 During pre-training,  $E_c$  and  $D_r$  are optimized  
 224 by minimizing an objective function  $\mathcal{L}_{pre}$  that  
 225 contains a reconstruction loss  $\mathcal{L}_{rec}$  and a latent  
 226 loss  $\mathcal{L}_{lat}$ , defined as follows:

$$\mathcal{L}_{pre} = \mathcal{L}_{rec} + \mathcal{L}_{lat} = \|I_f - I_r\|_1 + \alpha \|\text{sg}[Z_c] - Z_q\|_2^2 + \gamma \|Z_c - \text{sg}[Z_q]\|_2^2 \quad (2)$$

229 where  $\text{sg}$  denotes the stop-gradient operator.  $\alpha$  and  $\gamma$  are the balancing hyper-parameters. Experi-  
 230 mentally, we set them as 1 and 0.25, respectively.  
 231

232 Upon completing pre-training, we fix the content encoder  $E_c$  along with the codebook  $F_c$  to build  
 233 the font generation model. To be noticed, although the reconstruction decoder  $D_r$  adopts the same  
 234 architecture with the font generation model’s decoder, the latter is re-trained from scratch during the  
 235 main training phase.

236 **Content Alignment Module.** From a perceptual perspective, reference characters sharing similar  
 237 elements with the input glyph should receive more attention during the style transfer process (Zhu  
 238 et al., 2020). To this end, in the Content Alignment Module (CAM), we first extract the content  
 239 features  $f_s^c = \{f_i^c\}_{i=1}^k$  from reference images via the content encoder  $E_c$ . In addition, to ensure  
 240 dimension compatibility, both  $f_s^c$  and  $f_c$  are reshaped to  $\overline{f_s^c}$  and  $\overline{f_c}$ . Then, the similarity values are  
 241 calculated using the normalized cross-correlation measurement (Singhal, 2001) as:

$$\Phi_{ia} = \frac{\langle \overline{f_i^c}, \overline{f_c} \rangle}{\|\overline{f_i^c}\| * \|\overline{f_c}\|}, \quad a \in \{1, 2, \dots, T\} \quad (3)$$

245 where  $a$  refers to the position within the  $T$ -dimensional channel, and  $\Phi_{ia}$  is a scalar representing  
 246 the similarity between the  $a$ -th channel of the  $i$ -th reference image and the content image. Next,  
 247 we normalize these values channel-wise using the Softmax function, and apply them to weight the  
 248 spatial style feature representations  $\hat{f}_{spa}$  to obtain the aligned content feature  $f_{ca}$  by:

$$\overline{\Phi}_{ia} = \text{Softmax}(\Phi_{ia}) \quad f_{ca} = \text{Concat}_a \left( \sum_{i=1}^k \overline{\Phi}_{ia} \hat{f}_{is}^a \right) \quad (4)$$

### 253 3.4 STYLE-CONTENT FUSION MODULE

254 The Style-Content Fusion Module (SCFM) serves as the core integration unit that harmonizes con-  
 255 tent structure and style representations. As shown in Figure 4, each attention block contains two key  
 256 components, namely Grouped Residual Layer (GRL) and Exponential-Space Relative Position Bias  
 257 (ES-RPB), which are introduced below.

258 **Grouped Residual Layer.** Some prior works (Zhou et al., 2023; Li et al., 2024) have pointed out  
 259 the computational redundancy and complexity in the generation of query, key, and value matrices.  
 260 To mitigate this, we design the Grouped Residual Layer (GRL) that integrates grouping and residual  
 261 connections to improve efficiency and learning capacity. Formally, given an input  $S \in \mathbb{R}^{B \times N \times C}$ , we  
 262 first split  $S$  along the channel dimension into two halves, denoted as  $S_1, S_2 \in \mathbb{R}^{B \times N \times \frac{C}{2}}$ . Then, two  
 263 linear layers,  $\Gamma_1$  and  $\Gamma_2$ , are applied to both parts, combined with residual connections to enhance  
 264 the gradient flow, as follows:

$$\begin{aligned} \Omega_1 &= \Gamma_1(S_1) + S_1, \quad \Omega_1 \in \mathbb{R}^{B \times N \times \frac{C}{2}} \\ \Omega_2 &= \Gamma_2(S_2) + S_2, \quad \Omega_2 \in \mathbb{R}^{B \times N \times \frac{C}{2}} \end{aligned} \quad (5)$$

265 Ultimately,  $\Omega_1$  and  $\Omega_2$  are concatenated along the channel dimension to get the output.

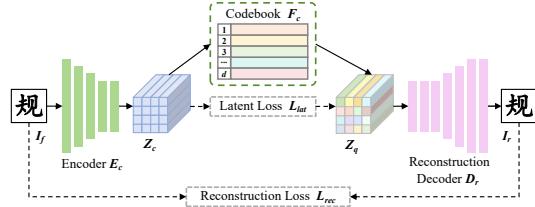


Figure 3: **The architecture of glyph feature decomposition network.** It is utilized for pre-training  $E_c$  and acquiring  $F_c$ .

270 **Exponential-Space Relative Position**  
 271 **Bias.** Considering that conventional Relative Position Bias (RPB) assigns positional  
 272 weights independently and overlooks the principle that nearby pixels contribute more to image generation (Liang  
 273 et al., 2021), we propose the Exponential-Space Relative Position Bias (ES-RPB).  
 274 Particularly, we introduce an exponential mapping to the original absolute position  
 275 coordinate ( $\Delta X, \Delta Y$ ), which injects distance-awareness into the original RPB,  
 276 thus promoting local focus. What's more, to better stabilize training, a lightweight  
 277 Multi-Layer Perceptron (MLP) is employed to process the transformed coordinate  
 278 ( $\Delta \hat{X}, \Delta \hat{Y}$ ) and acquire the final bias matrix  $\Lambda_E$ . The operation process can be  
 279 formulated as:  
 280

$$\begin{aligned} \Delta \hat{X} &= \text{sign}(\Delta X) \cdot (1 - \exp(-|\delta \cdot \Delta X|)) \\ \Delta \hat{Y} &= \text{sign}(\Delta Y) \cdot (1 - \exp(-|\eta \cdot \Delta Y|)) \\ \Lambda_E &= \text{MLP}(\Delta \hat{X}, \Delta \hat{Y}) \end{aligned} \quad (6)$$

281 where  $\delta$  and  $\eta$  are the trainable factors that control the sensitivity to positional offsets within the same  
 282 window. Besides, the symbol  $\cdot$  denotes element-wise multiplication, and the MLP comprises two  
 283 linear layers with an activation function in between. We set both  $\delta$  and  $\eta$  to 3.2 in our experiments.  
 284

285 Based on the GRL and ES-RPB, we elaborate on the fusion process of SCFM. Notably, to ensure dimension  
 286 matching, we first flatten the spatial structure of key and value features across both blocks.  
 287 In each GRAB, the query  $Q$ , key  $K$ , and value  $V$  are derived by applying three GRLs to the corresponding  
 288 inputs. Next, the similarity between  $Q$  and  $K$  is computed via normalized dot-product, which is then scaled by a learnable factor  $\mu$  (initialized to 10), and further refined with  $\Lambda_E$ . After that, the attention weights are obtained through the Softmax function, followed by feature aggregation with  $V$ . Lastly, a linear projection  $\Gamma_{Proj}$  is employed to output the result. The overall pipeline is represented as follows:  
 289

$$\text{GRAB} = \Gamma_{proj}(\text{Softmax}(\mu \cdot \text{Normalize}(Q) * \text{Normalize}(K)^\top + \Lambda_E)V) \quad (7)$$

303 In the first GRAB, the component-wise codebook  $F_c$  serves as  $Q$ , while the refined spatial style  
 304 feature  $\hat{f}_{spa}$  is used as  $K$  and  $V$ , yielding the intermediate output  $F_s$ . The second GRAB follows  
 305 the same structure, where  $Q$  is replaced by  $F_s$ , with both  $K$  and  $V$  jointly assigned to  $f_{fre}$ . Finally,  
 306 we utilize an output adapter to adjust the channel dimension, obtaining the fused style representation  
 307  $f_{sa}$ . More theoretical analyses of SCFM can be found in Appendix A.5.  
 308

### 3.5 TRAINING OBJECTIVE

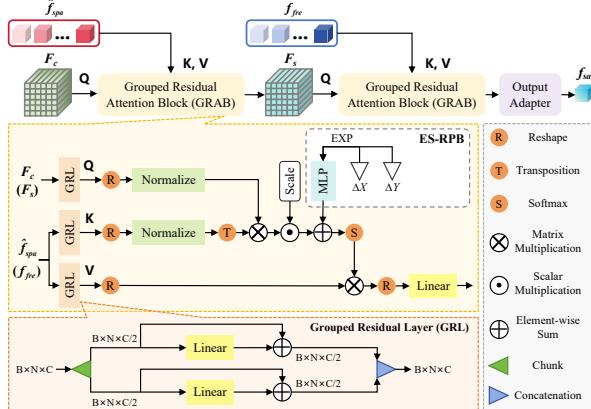
313 The loss functions of our proposed HF-Font consist of three parts: adversarial loss, matching loss,  
 314 and style contrastive loss.

315 **Adversarial Loss.** To ensure plausibility in both style and content, we use a multi-head discriminator  
 316  $D_{s,c}$  conditioned on the style label  $s$  and content label  $c$ . The loss function is implemented on  
 317 the hinge GAN loss (Zhang et al., 2019) by:

$$\begin{aligned} \mathcal{L}_{adv}^D &= -\mathbb{E}_{I_t \sim p_{data}} \min(0, -1 + D_{s,c}(I_t)) - \mathbb{E}_{I_o \sim p_G} \min(0, -1 - D_{s,c}(I_o)) \\ \mathcal{L}_{adv}^G &= -\mathbb{E}_{I_o \sim p_G} D_{s,c}(I_o) \end{aligned} \quad (8)$$

321 where  $p_{data}$  and  $p_G$  denote the set of real images and generated images, respectively.

322 **Matching Loss.** To mitigate mode collapse and enforce the generated character  $I_o$  closely resemble  
 323 to the ground truth  $I_t$  at both pixel and feature levels, we apply an  $\mathcal{L}_1$  loss on the image and its



324 **Figure 4: Style-Content Fusion Module.** It is made up  
 325 of two Grouped Residual Attention Blocks (GRABs)  
 326 and one output adapter for dimensional adjustment.

324 features, as defined below:

$$\begin{aligned} \mathcal{L}_{img} &= \mathbb{E}_{I_t \sim p_{data}} [\|I_t - I_o\|_1] \\ \mathcal{L}_{feat} &= \mathbb{E}_{I_t \sim p_{data}} \left[ \sum_{m=1}^M \left\| D_{s,c}^{(m)}(I_t) - D_{s,c}^{(m)}(I_o) \right\|_1 \right] \end{aligned} \quad (9)$$

330 where  $M$  is  $D_{s,c}$ ’s layer number and  $D_{s,c}^{(m)}(\cdot)$  refers to the feature map in the  $m$ -th layer of  $D_{s,c}$ .

331 **Style Contrastive Loss.** To promote the extraction of discriminative style features from high-  
332 frequency information, we propose the style contrastive loss  $\mathcal{L}_{cst}$ . It encourages high-frequency  
333 style feature  $f_{fre}$  from the same font to be closer in the feature space, while pushing apart those  
334 from different fonts, as follows:

$$\mathcal{L}_{cst} = -\frac{1}{U} \sum_{i \in \Theta} \frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(v_i \cdot v_p / \tau)}{\sum_{a \in A(i)} \exp(v_i \cdot v_a / \tau)} \quad (10)$$

339 In detail,  $i \in \Theta = \{1, 2, \dots, U\}$  is the sample index in a batch size of  $U$  and  $A(i) = \Theta \setminus \{i\}$  represents  
340 other indices distinct from  $i$ . Besides,  $v_i$  denotes the normalized high-frequency feature of the  $i$ -th  
341 font derived from  $f_{fre}^{(i)}$ . Its positive set is defined as  $P(i) = \{p \in A(i) \mid y_p = y_i\}$ , where  $y_i$  refers  
342 to the font label of the  $i$ -th sample, and the remaining samples  $A(i) \setminus P(i)$  serve as the negative  
343 set. Apart from these, the scalar parameter  $\tau$  controls the temperature of the similarity distribution,  
344 which is set to 0.07.

345 **Total Loss.** Finally, we optimize HF-Font by the following full objective function:

$$\min_G \max_D (\mathcal{L}_{adv}^D + \mathcal{L}_{adv}^G + \lambda_1 \mathcal{L}_{img} + \lambda_2 \mathcal{L}_{feat} + \lambda_3 \mathcal{L}_{cst}) \quad (11)$$

349 Here,  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are the three weighting hyper-parameters. In our experiments, we empirically  
350 set them to 1, 1, and 0.2, respectively.

## 352 4 EXPERIMENTS

### 354 4.1 DATASET AND EVALUATION METRICS

356 **Dataset.** We collect a large Chinese font dataset with 575 fonts (style), each containing 3500  
357 commonly-used Chinese characters (content) at a resolution of  $128 \times 128$ . Notably, the font *kai*  
358 is fixed as the content font throughout training and testing. It is also used to pre-train the feature  
359 decomposition network for codebook acquisition.

360 For the training set, we randomly select 550 fonts with 3000 characters per font, forming the Seen  
361 Fonts Seen Characters (SFSC) set. Our test set includes two parts: 24 Unseen Fonts with 500 Unseen  
362 Characters per font (UFUC) and 550 Seen Fonts with 500 Unseen Characters per font (SFUC).

363 **Evaluation Metrics.** To comprehensively assess the quality of font generation, we utilize the  
364 following five metrics, including  $L_1$  loss, Structural Similarity Index Measure (SSIM) (Wang et al.,  
365 2004), Root Mean Square Error (RMSE) (Karunasingha, 2022), Learned Perceptual Image Patch  
366 Similarity (LPIPS) (Zhang et al., 2018a), and Frechet Inception Distance (FID) (Martin et al., 2017).

367 Furthermore, we also perform a user study to evaluate the subjective quality of generated results. In  
368 particular, we randomly select 10 font styles from each test set and 10 characters per font. Then,  
369 we invite 20 well-educated volunteers and ask them to choose the best generation result among the  
370 comparison methods. Here, all the samples are shuffled before evaluation to avoid any potential  
371 bias. **In addition, all the participants evaluate the same set of samples to ensure consistency and  
372 comparability across the results.**

### 374 4.2 IMPLEMENTATION DETAILS

376 The entire training process consists of two stages. In the first stage, we train the feature decomposing  
377 network with 3000 Chinese characters rendered in the font *kai*. During this stage, we set the em-  
bedding dimension to 256, the codebook size to 100, the batch size to 64, and the number of training

378 Table 1: Quantitative comparison results on UFUC and SFUC datasets.  
379

380 Dataset	381 Method	382 Venue	383 SSIM↑	384 RMSE↓	385 LPIPS↓	386 FID↓	387 $L_1 \downarrow$	388 User study↑
389 UFUC	FUNIT	ICCV 2019	0.6276	0.3356	0.2692	71.3708	0.1376	2.55%
	MX-Font	ICCV 2021	0.6966	0.3159	0.2359	62.7500	0.1235	5.25%
	DG-Font	CVPR 2021	0.6527	0.3238	0.2058	61.6022	0.1241	5.55%
	LF-Font	AAAI 2021	0.6768	0.3110	0.2516	66.8840	0.1190	5.45%
	CF-Font	CVPR 2023	0.6613	0.3102	0.2014	59.7644	0.1169	8.90%
	VQ-Font	ICCV 2023	0.6776	0.3066	0.2157	58.2632	0.1175	10.35%
	FontDiffuser	AAAI 2024	0.6390	0.3579	0.2816	79.4458	0.1530	8.35%
	IF-Font	NeurIPS 2024	0.6891	0.3130	0.2173	59.6292	0.1038	11.10%
	MX-Font++	ICASSP 2025	<b>0.7092</b>	<b>0.2897</b>	<b>0.1855</b>	<b>57.7845</b>	<b>0.0976</b>	<b>12.05%</b>
	DA-Font	ACM MM 2025	<b>0.7352</b>	<b>0.2854</b>	<b>0.1699</b>	<b>54.1856</b>	<b>0.0872</b>	<b>13.75%</b>
390 SFUC	HF-Font (Ours)	-	<b>0.7518</b>	<b>0.2692</b>	<b>0.1544</b>	<b>50.0632</b>	<b>0.0817</b>	<b>16.70%</b>
	FUNIT	ICCV 2019	0.6056	0.3583	0.2777	68.8726	0.1545	2.25%
	MX-Font	ICCV 2021	0.6733	0.3314	0.2197	61.0583	0.1321	5.65%
	DG-Font	CVPR 2021	0.6478	0.3278	0.2006	58.3525	0.1277	5.80%
	LF-Font	AAAI 2021	0.6140	0.3335	0.2726	69.2239	0.1520	5.55%
	CF-Font	CVPR 2023	0.6569	0.3186	0.1974	57.9868	0.1202	8.60%
	VQ-Font	ICCV 2023	0.6414	0.3304	0.2069	56.7299	0.1305	9.95%
	FontDiffuser	AAAI 2024	0.6317	0.3691	0.2910	73.3052	0.1574	8.65%
	IF-Font	NeurIPS 2024	0.6652	0.3280	0.2221	62.8246	0.1213	10.70%
	MX-Font++	ICASSP 2025	<b>0.6843</b>	<b>0.3122</b>	<b>0.2043</b>	<b>57.1191</b>	<b>0.1179</b>	<b>11.40%</b>
391	DA-Font	ACM MM 2025	<b>0.7287</b>	<b>0.3019</b>	<b>0.1792</b>	<b>49.2237</b>	<b>0.1116</b>	<b>14.05%</b>
	HF-Font (Ours)	-	<b>0.7335</b>	<b>0.2941</b>	<b>0.1677</b>	<b>46.1692</b>	<b>0.1063</b>	<b>17.40%</b>

392 Dataset	393 UFUC					394 SFUC				
395 Content	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
FUNIT	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
MX-Font	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
DG-Font	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
LF-Font	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
CF-Font	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
VQ-Font	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
FontDiffuser	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
IF-Font	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
MX-Font++	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
DA-Font	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
HF-Font (Ours)	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快
Ground Truth	激 货 草 辉	鲸 架 芥 焦	轨 故 拱 给	荒 戈 哒 虹	桂 管 耿 港	旧 晶 监 寄	卷 驹 崛 钩	康 科 垦 快	康 科 垦 快	康 科 垦 快

411 Figure 5: Qualitative comparison results on UFUC and SFUC datasets. Characters in the blue  
412 boxes suffer from stroke errors (including missing or redundant strokes), while the red boxes high-  
413 light conspicuous blurring or artifacts. Zoom in to see details.  
414415 iterations to 50000. In the second stage, we train the entire model via Adam optimizer (Kingma &  
416 Ba, 2015) with  $\beta_1 = 0.9$  and  $\beta_2 = 0.99$ . Here, the batch size is set to 8, with learning rates of  
417  $2 \times 10^{-4}$  for the generator and  $4 \times 10^{-4}$  for the discriminator. The total number of iteration steps is  
418 set to 600000. For the few-shot font generation task, the number of reference images is set to 4.  
419420 4.3 COMPARISON WITH SOTA METHODS  
421422 We compare our proposed HF-Font with eight state-of-the-art methods, including one image-to-  
423 image translation method (FUNIT (Liu et al., 2019)) and nine classical font generation methods  
424 (MX-Font (Park et al., 2021a), DG-Font (Xie et al., 2021), LF-Font (Park et al., 2021b), CF-  
425 Font (Wang et al., 2023), VQ-Font (Pan et al., 2023), FontDiffuser (Yang et al., 2024), IF-Font (Chen  
426 et al., 2024), MX-Font++ (Wang et al., 2025), and DA-Font (Chen et al., 2025)). For a fair compar-  
427 ison, we retrain all these models using their default settings on our training set. To be noticed, none  
428 of these above methods explicitly leverage high-frequency information.  
429430 The quantitative comparison results are summarized in Table 1. It is evident that our proposed  
431 model achieves the best performance across all the evaluation metrics. Figure 5 displays the qual-  
432 itative comparison results. We can observe that FUNIT exhibits structural incompleteness in most

432  
433 **Table 2: Quantitative results on different**  
434 **modules.** Among them, G, H, F, and L denote  
435 the gate mechanism, high-frequency branch,  
SCFM, and style contrastive loss, respectively.

437	Modules	G	H	F	L	SSIM↑	RMSE↓	LPIPS↓	FID↓
438						0.6434	0.3228	0.2481	68.7447
439	✓					0.6560	0.3175	0.2293	63.9206
440	✓	✓				0.6953	0.2986	0.1985	59.1434
441	✓	✓	✓			0.7367	0.2734	0.1679	54.2591
442	✓	✓	✓	✓		<b>0.7518</b>	<b>0.2692</b>	<b>0.1544</b>	<b>50.0632</b>

443	Content	借	鉴	绩	戒	駿	咖	军	救	尽	节	江	进	炕	肯	恐	库
444	Base Model	借	鉴	绩	戒	駿	咖	军	救	尽	节	江	进	炕	肯	恐	库
445	+G	借	鉴	绩	戒	駿	咖	军	救	尽	节	江	进	炕	肯	恐	库
446	+GH	借	鉴	绩	戒	駿	咖	军	救	尽	节	江	进	炕	肯	恐	库
447	+GHF	借	鉴	绩	戒	駿	咖	军	救	尽	节	江	进	炕	肯	恐	库
448	+GHFL	借	鉴	绩	戒	駿	咖	军	救	尽	节	江	进	炕	肯	恐	库
449	Ground Truth	借	鉴	绩	戒	駿	咖	军	救	尽	节	江	进	炕	肯	恐	库

450 **Figure 6: Qualitative results on different modules.** G, H, F, and L share the same notations  
451 with Table 2. The first row is the base model.  
452 The green boxes highlight details better generated  
453 by the full model. Zoom in to see details.

454  
455  
456 cases. Only when the source font closely resembles the target one can it generate well-layout char-  
457 acters. Although MX-Font and LF-Font could preserve the overall character shapes, their outputs  
458 often appear blurry with indistinct textures, resulting in degraded stroke smoothness and clarity.  
459 Besides, DG-Font also struggles to capture fine-grained details in complex glyphs, which leads to  
460 incomplete or fuzzy components. While CF-Font generally maintains the correct glyph layouts,  
461 it usually introduces noticeable artifacts. VQ-Font demonstrates good stability but lacks precise  
462 control over local features, resulting in glyph distortions and vague strokes. In addition, the per-  
463 formances of FontDiffuser and IF-Font are also outstanding, but their generated results sometimes  
464 suffer from obvious stroke errors. [MX-Font++ improves style aggregation yet still shows obvious](#)  
465 [drawbacks in some cases, while DA-Font could preserve well structures but sometimes suffers from](#)  
466 [local blurring.](#) As a whole, characters generated by our HF-Font are of high quality in terms of style  
467 consistency, accurate global structures, and delicate local details. [Notably, Additional comparison](#)  
468 [results are provided in Appendix A.3, more generation results are included in Appendix A.6, and](#)  
469 [the cross-lingual generation results are presented in Appendix A.7.](#)

#### 470 4.4 ABLATION STUDY

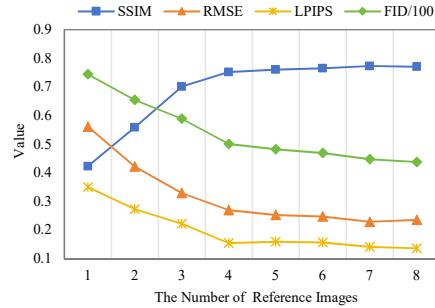
471 In this section, we conduct a series of ablation experiments to access the influence of each part in  
472 our model. These experiments are performed on the UFUC dataset. More analyses and results are  
473 provided in the Appendix A.4.

474 **The Effect of Different Modules.** We separate the proposed modules and sequentially add them  
475 to the base model to observe their individual effects. The quantitative results are presented in Ta-  
476 ble 2, validating that all these modules can help improve the quality. Beyond numerical advances,  
477 they also bring a noticeable improvement in the visual aspects of geometric structures and stylistic  
478 strokes, as depicted in Figure 6. The gate mechanism suppresses background noise to enable more  
479 reliable spatial style extraction. Building on this, the high-frequency branch enriches local details  
480 and SCFM facilitates comprehensive feature fusion, which strengthens both structural coherence  
481 and style consistency. Moreover, the style contrastive loss further enhances the precision of results.

482 **The Effect of Codebook Size.** The size of codebook  $F_c$  significantly affects the model’s complexity  
483 and feature decomposition ability. Table 3 presents the quantitative results under different codebook  
484 sizes. We can find that the performance improves with increasing codebook size up to 100. Never-

485 **Table 3: Quantitative results on different**  
486 **codebook sizes.** It demonstrates how varying  
487 the number of embeddings  $F_c$  affects the qual-  
488 ity of generated images.

489	Codebook Size	SSIM↑	RMSE↓	LPIPS↓	FID↓
490	50	0.6276	0.3675	0.2531	69.6119
491	75	0.7173	0.3257	0.2074	62.0768
492	100	<b>0.7518</b>	<b>0.2692</b>	<b>0.1544</b>	<b>50.0632</b>
493	125	0.7624	0.2704	0.1456	49.1751
494	150	0.7689	0.2698	0.1497	47.0893



495 **Figure 7: Performance trend with different**  
496 **numbers of reference images.**

486  
 487   theless, beyond this point, further enlarging the codebook yields marginal gains, with some metrics  
 488   (e.g., RMSE) even showing slight degradation. This suggests that while a larger codebook could  
 489   enhance representation, excessive sizes would introduce more redundant information, which might  
 490   compromise the model’s efficiency. Hence, in other experiments, we set the codebook size to 100.

491   **The Effect of Reference Image Numbers.** Intuitively, given more reference images, the generated  
 492   results would be better due to the availability of richer style information. As illustrated in Fig-  
 493   ure 7, the model’s performance consistently improves as the number of reference images increases.  
 494   Specifically, when the number rises from 1 to 4, the quality of generated results yields notable gains.  
 495   However, further increases bring only marginal improvements while introducing additional infe-  
 496   rference time and memory overhead. Thus, in our paper, to strike a balance between performance and  
 497   efficiency, we set the number of reference images to 4.

## 498   5 CONCLUSION

500   In this paper, we propose HF-Font, a novel few-shot font generation framework that leverages high-  
 501   frequency information from reference images to enhance style extraction. To capture more distinc-  
 502   tive style features, we introduce a style contrastive loss to facilitate the transfer process. Meanwhile,  
 503   a gate mechanism is employed to suppress background noise while retaining meaningful style cues.  
 504   Furthermore, we also design a Style-Content Fusion Module (SCFM) to promote effective inte-  
 505   gration of style and content representations. Both quantitative and qualitative experimental results  
 506   verify that our proposed model exceeds other competitive methods on various fonts and characters.

## 508   REPRODUCIBILITY STATEMENT

510   We have included detailed descriptions of our method in the main text (Section 3), as well as exper-  
 511   imental settings (Section 4.1 and Section 4.2) to ensure reproducibility. Additional information is  
 512   provided in the appendix, including the network architecture (Appendix A.1) and theoretical analy-  
 513   ses (Appendix A.5), which further support the reproducibility of our results.

## 515   REFERENCES

517   van den Oord Aaron, Vinyals Oriol, Kavukcuoglu Koray, and et al. Neural discrete representation  
 518   learning. In *Proc. NeurIPS*, pp. 6309–6318, 2017.

519   Weiran Chen, Chunping Liu, Yi Ji, and et al. Chinese character style transfer model based on  
 520   convolutional neural network. In *Proc. ICANN*, pp. 558–569, 2022.

522   Weiran Chen, Guiqian Zhu, Ying Li, Yi Ji, and Chunping Liu. DA-Font: few-shot font generation  
 523   via dual-attention hybrid integration. In *Proc. ACM MM*, pp. 6644–6653, 2025.

524   Xinping Chen, Xiao Ke, Wenzhong Guo, and et al. IF-Font: Ideographic description sequence-  
 525   following font generation. In *Proc. NeurIPS*, pp. 1–23, 2024.

527   Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, and et al. ILVR: Conditioning method for denois-  
 528   ing diffusion probabilistic models. In *Proc. ICCV*, pp. 14347–14356, 2021.

529   Gang Dai, Yifan Zhang, Quhui Ke, and et al. One-DM: One-shot diffusion mimicker for handwritten  
 530   text generation. In *Proc. ECCV*, pp. 410–427, 2024.

531   Bin Fu, Junjun He, Jianjun Wang, and Yu Qiao. Neural transformation fields for arbitrary-styled  
 532   font generation. In *Proc. CVPR*, pp. 22438–22447, 2023.

534   Yue Gao, Yuan Guo, Zhouhui Lian, and et al. Artistic glyph image synthesis via one-stage few-shot  
 535   learning. *ACM Trans. Graph.*, 38:1–12, 2020.

536   Ian J Goodfellow, Pouget-Abadie Jean, Mirza Mehdi, Bing Xu, Aaron Courville, and et al. Genera-  
 537   tive adversarial nets. In *Proc. NeurIPS*, pp. 2672–2680, 2014.

539   Haibing He, Xinyuan Chen, Chaoyue Wang, and et al. Diff-Font: Diffusion model for robust one-  
 540   shot font generation. *Int. J. Comput. Vis.*, 132:5372–5386, 2024.

540 Jonathan Ho, Ajay Jain, Pieter Abbeel, and et al. Denoising diffusion probabilistic models. In *Proc. NeurIPS*, pp. 1–12, 2020.

541

542

543 Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proc. ICCV*, pp. 1510–1519, 2017.

544

545 Philip Isola, Jun-Yan Zhu, Tinghui Zhou, and et al. Image-to-image translation with conditional adversarial networks. In *Proc. CVPR*, pp. 5967–5976, 2017.

546

547

548 Dulakshi Santhusitha Kumari Karunasingha. Root mean square error or mean absolute error? use their ratio as well. *Inf. Sci.*, 585:609–629, 2022.

549

550 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. ICLR*, pp. 1–15, 2015.

551

552

553 Juyeon Ko, Inho Kong, Dogyun Park, and et al. Stochastic conditional diffusion models for robust semantic image synthesis. In *Proc. ICML*, pp. 24932–24963, 2024.

554

555 Yuxin Kong, Canjie Luo, Weihong Ma, Qiyuan Zhu, Shenggao Zhu, Nicholas Yuan, and et al. Look closer to supervise better: One-shot font generation via component-based discriminator. In *Proc. CVPR*, pp. 13472–13481, 2022.

556

557

558 Yuzhen Li, Zehang Deng, Yuxin Cao, and et al. GRFormer: Grouped residual self-attention for lightweight single image super-resolution. In *Proc. ACM MM*, pp. 9378–9386, 2024.

559

560

561 Jingyun Liang, Jiazhang Cao, Guolei Sun, Kai Zhang, Radu Timofte, and et al. SwinIR: Image restoration using swin transformer. In *Proc. ICCV*, pp. 1833–1844, 2021.

562

563

564 Mingyu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, and et al. Few-shot unsupervised image-to-image translation. In *Proc. ICCV*, pp. 10550–10559, 2019.

565

566 Heusel Martin, Ramsauer Hubert, Unterthiner Thomas, and et al. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proc. NeurIPS*, pp. 6629–6640, 2017.

567

568

569 Ben Mildenhall, Pratul P. Srinivasan, and et al. Nerf: Representing scenes as neural radiance fields for view synthesis. In *Proc. ECCV*, pp. 405–421, 2020.

570

571 Wei Pan, Anna Zhu, Xinyu Zhou, and et al. Few shot font generation via transferring similarity guided global style and quantization local style. In *Proc. ICCV*, pp. 19449–19459, 2023.

572

573

574 Song Park, Sanghyuk Chun, Junbum Cha, and et al. Multiple heads are better than one: Few-shot font generation with multiple localized experts. In *Proc. ICCV*, pp. 13880–13889, 2021a.

575

576 Song Park, Sanghyuk Chun, Junbum Cha, and et al. Few-shot font generation with localized style representations and factorization. In *Proc. AAAI*, pp. 2393–2402, 2021b.

577

578

579 Amit Singhal. Modern information retrieval: A brief overview. *IEEE Data Eng. Bull.*, 24:35–43, 2001.

580

581 Danyang Sun, Tongzheng Ren, Chongxuan Li, and et al. Learning to write stylized chinese characters by reading a handful of examples. In *Proc. IJCAI*, pp. 920–927, 2018.

582

583

584 Licheng Tang, Yiyang Cai, Jiaming Liu, Zhibin Hong, and et al. Few-shot font generation by learning fine-grained local styles. In *Proc. CVPR*, pp. 7895–7904, 2022.

585

586 Martin Wahl. A kernel-based analysis of laplacian eigenmaps, 2024.

587

588 Chi Wang, Min Zhou, Ge Tiezheng, and et al. CF-Font: Content fusion for few-shot font generation. In *Proc. CVPR*, pp. 1858–1867, 2023.

589

590 Weihang Wang, Duolin Sun, Jielei Zhang, and Longwen Gao. MX-Font++: mixture of heterogeneous aggregation experts for few-shot font generation. In *Proc. ICASSP*, pp. 1–5, 2025.

591

592

593 Zhou Wang, A.C Bovik, H.R. Sheikh, and et al. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13:600–612, 2004.

594 Yangchen Xie, Xinyuan Chen, Li Sun, and Yue Lu. DG-Font: Deformable generative networks for  
595 unsupervised font generation. In *Proc. CVPR*, pp. 5130–5140, 2021.

596  
597 Zhenhua Yang, Dezhi Peng, Yuxin Kong, and et al. FontDiffuser: one-shot font generation via  
598 denoising diffusion with multi-scale content aggregation and style contrastive learning. In *Proc.*  
599 *AAAI*, pp. 6603–6611, 2024.

600 Han Zhang, Ian J. Goodfellow, Dimitris N. Metaxas, and et al. Self-attention generative adversarial  
601 networks. In *Proc. ICML*, pp. 7354–7363, 2019.

602 Richard Zhang, Phillip Isola, A.Efros Alexei, and et al. The unreasonable effectiveness of deep  
603 features as a perceptual metric. In *Proc. CVPR*, pp. 586–595, 2018a.

604  
605 Yexun Zhang, Ya Zhang, Wenbin Cai, and et al. Separating style and content for generalized style  
606 transfer. In *Proc. CVPR*, pp. 8447–8455, 2018b.

607  
608 Yupeng Zhou, Zhen Li, Chunle Guo, and et al. SRFormer: Permuted self-attention for single image  
609 super-resolution. In *Proc. ICCV*, pp. 12734–12745, 2023.

610 Anna Zhu, Xiongbo Lu, Xiang Bai, and et al. Few-shot text style transfer via deep feature similarity.  
611 *IEEE Trans. Image Process.*, 29:6932–6946, 2020.

612 Jun-Yan Zhu, Taesung Park, Isola Philip, and et al. Unpaired image-to-image translation using  
613 cycle-consistent adversarial networks. In *Proc. ICCV*, pp. 2242–2251, 2017.

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

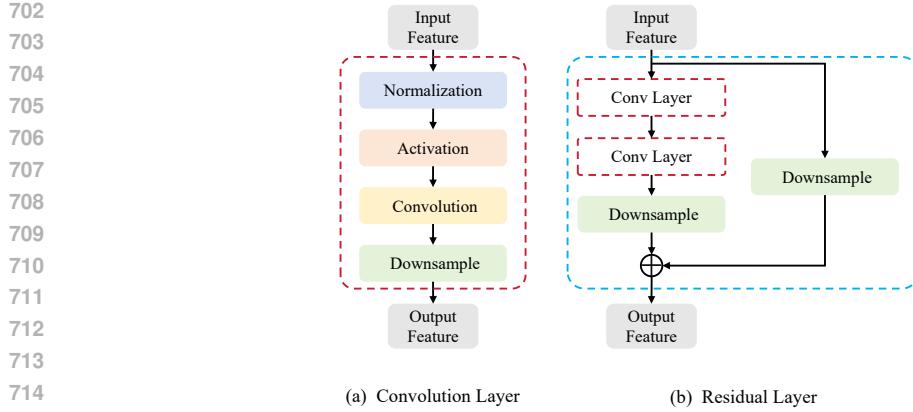
645

646

647

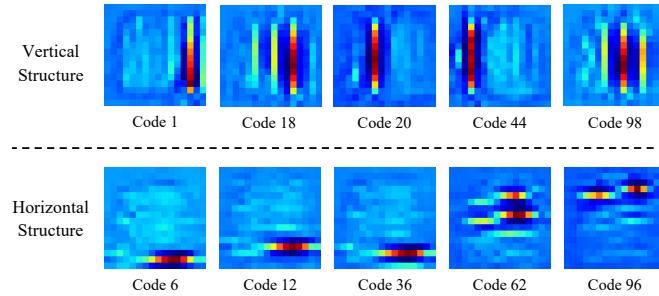
648 **A APPENDIX**  
649650 In this appendix, we provide additional details and analyses, including network architecture, visu-  
651 alizations of the codebook, extended comparison and ablation studies, theoretical analyses of GRL  
652 and ES-RPB in SCFM, more generation results, [cross-lingual generation results](#), as well as a brief  
653 description of our usage of large language models (LLMs) for writing polish.  
654655 **A.1 NETWORK ARCHITECTURE**  
656657 Our entire model consists of two primary parts: the generator and the discriminator. Both of them  
658 are built up with two types of layers: the convolution layer and the residual layer, which are illus-  
659 trated in Figure 8. Each residual layer includes two identical convolution layers. Notably, since  
660 the residual layer has its own downsampling operator, its convolution layers do not perform addi-  
661 tional downsampling. The detailed architectures of content encoder  $E_c$ , spatial style encoder  $E_{ss}$ ,  
662 high-freq style encoder  $E_{sh}$ , decoder  $D_r$ , and discriminator  $D_{s,c}$  are shown in Table 4.  
663664 **Table 4: Architectures of the generator modules  $E_c$ ,  $E_{ss}$ ,  $E_{sh}$ ,  $D_r$ , and the discriminator  $D_{s,c}$ .**  
665 IN and SN denote instance normalization and spectral normalization, respectively. Besides, all the  
666 padding operations are zero-padding. In the discriminator’s output layer, we utilize two embedding  
667 operators to embed the output feature map into two prediction vectors of both its style and content.  
668

Content Encoder $E_c$								
Layer Type	Normalization	Activation	Padding	Kernel Size	Stride	Downsample	Feature Maps	
Convolution Layer	IN	ReLU	1	3	1	-	32	
Convolution Layer	IN	ReLU	1	3	2	-	64	
Convolution Layer	IN	ReLU	1	3	2	-	128	
Convolution Layer	IN	ReLU	1	3	2	-	256	
Convolution Layer	IN	ReLU	1	3	1	-	256	
Spatial Style Encoder $E_{ss}$								
Layer Type	Normalization	Activation	Padding	Kernel Size	Stride	Downsample	Feature Maps	
Convolution Layer	IN	ReLU	1	3	1	-	32	
Convolution Layer	IN	ReLU	1	3	1	AvgPool	64	
Convolution Layer	IN	ReLU	1	3	1	AvgPool	128	
Residual Layer	IN	ReLU	1	3	1	-	128	
Residual Layer	IN	ReLU	1	3	1	-	128	
Residual Layer	IN	ReLU	1	3	1	AvgPool	256	
Residual Layer	IN	ReLU	1	3	1	-	256	
Output Layer	-	Sigmoid	-	-	-	-	256	
High-freq Style Encoder $E_{sh}$								
Layer Type	Normalization	Activation	Padding	Kernel Size	Stride	Downsample	Feature Maps	
Convolution Layer	IN	ReLU	1	3	1	-	32	
Convolution Layer	IN	ReLU	1	3	1	AvgPool	64	
Convolution Layer	IN	ReLU	1	3	1	AvgPool	128	
Residual Layer	IN	ReLU	1	3	1	-	128	
Residual Layer	IN	ReLU	1	3	1	-	128	
Residual Layer	IN	ReLU	1	3	1	AvgPool	256	
Residual Layer	IN	ReLU	1	3	1	-	256	
Output Layer	-	Sigmoid	-	-	-	-	256	
Decoder $D_r$								
Layer Type	Normalization	Activation	Padding	Kernel Size	Stride	Upsample	Feature Maps	
Residual Layer	IN	ReLU	1	3	1	-	256	
Residual Layer	IN	ReLU	1	3	1	-	256	
Residual Layer	IN	ReLU	1	3	1	-	256	
Convolution Layer	IN	ReLU	1	3	1	Nearest	128	
Convolution Layer	IN	ReLU	1	3	1	Nearest	64	
Convolution Layer	IN	ReLU	1	3	1	Nearest	32	
Convolution Layer	IN	ReLU	1	3	1	-	1	
Output Layer	-	Sigmoid	-	-	-	-	1	
Discriminator $D_{s,c}$								
Layer Type	Normalization	Activation	Padding	Kernel Size	Stride	Upsample	Feature Maps	
Convolution Layer	SN	-	1	3	2	-	32	
Residual Layer	SN	ReLU	1	3	1	AvgPool	64	
Residual Layer	SN	ReLU	1	3	1	AvgPool	128	
Residual Layer	SN	ReLU	1	3	1	AvgPool	256	
Residual Layer	SN	ReLU	1	3	1	-	512	
Residual Layer	SN	ReLU	1	3	1	-	512	
Output Layer						Multi-task embedding		

Figure 8: **Structures of the convolution layer and the residual layer in Table 4.**

## A.2 VISUAL ANALYSES OF THE CODEBOOK

To gain a deeper understanding into the learned quantization space, we visualize some representative codebook vectors to illustrate the structural information they encode. As shown in Figure 9, each code vector exhibits distinct spatial activations, with some focusing on vertical structures and others highlighting horizontal patterns. This indicates that the codebook could implicitly capture component-level abstractions, which facilitates structural decomposition without the need for any explicit prior knowledge. Such findings provide empirical support for the interpretability of the learned codebook and its critical role in character layout modelling.

Figure 9: **The visualization of some latent codes.**

## A.3 MORE COMPARISON RESULTS

In addition to the two test sets (UFUC and SFUC) discussed in the main paper, we further conduct comparison experiments on a third benchmark, 24 Unseen Fonts with 3000 Seen Characters per font (UFSC), to evaluate the generalization ability of our proposed method.

Apart from the state-of-the-art methods mentioned in Section 4.3, we also include Diff-Font, a generative architecture built upon the diffusion model, as an extra comparison method. Notably, since Diff-Font is limited to generating seen characters, we only evaluate it on the UFSC test set.

The quantitative comparison results are presented in Table 5, which shows that our HF-Font consistently outperforms all the competing methods across multiple evaluation metrics. Figure 10 displays the qualitative comparison results. We can observe that FUNIT tends to produce structurally incomplete glyphs, while MX-Font and LF-Font preserve the overall shapes but yield blurred textures with limited visual fidelity. DG-Font struggles with local details, leading to missing strokes or misplacing components in complex characters. Besides, CF-Font maintains the overall layout but suffers from artifacts and style inconsistencies. VQ-Font achieves reasonable style adaptation yet lacks fine-grained control, resulting in glyph distortions. Although Diff-Font, FontDiffuser, and IF-Font perform relatively well, stroke-level errors still occur in some cases. **MX-Font++ enhances style ag-**

756 gregation but still exhibits noticeable issues in certain cases, whereas DA-Font generally preserves  
 757 global structures yet occasionally suffers from local blurriness. In contrast, our HF-Font generates  
 758 characters with superior style consistency, better structural accuracy, and refined local details.  
 759

760 Table 5: Quantitative comparison results on UFSC dataset.  
 761

Method	Venue	SSIM↑	RMSE↓	LPIPS↓	FID↓	L1↓	User study↑
FUNIT	ICCV 2019	0.6309	0.3339	0.2685	59.3728	0.1362	2.40%
MX-Font	ICCV 2021	0.7062	0.3229	0.2063	53.3935	0.1205	4.30%
DG-Font	CVPR 2021	0.6658	0.3138	0.1877	50.0568	0.1256	4.70%
LF-Font	AAAI 2021	0.6812	0.3090	0.2471	55.2930	0.1176	4.60%
CF-Font	CVPR 2023	0.6794	0.3063	0.1836	49.7423	0.1224	7.95%
VQ-Font	ICCV 2023	0.6842	0.2995	0.1927	48.6191	0.1166	9.75%
Diff-Font <sup>1</sup>	IJCV 2024	0.6143	0.3242	0.1851	60.0310	0.1370	6.85%
FontDiffuser	AAAI 2024	0.6430	0.3433	0.2846	63.1833	0.1483	8.20%
IF-Font	NeurIPS 2024	0.6915	0.2997	0.1962	48.3104	0.1181	10.65%
MX-Font++	ICASSP 2025	0.7351	0.2923	0.1798	46.7172	0.1127	11.55%
DA-Font	ACM MM 2025	0.7544	0.2782	0.1649	43.8763	0.0916	13.60%
HF-Font (Ours)	-	0.7612	0.2635	0.1567	40.6694	0.0859	15.45%

771 <sup>1</sup> Diff-Font can only generate seen characters, so we evaluate this model on the UFSC test set alone.  
 772

Content	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
FUNIT	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
MX-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
DG-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
LF-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
CF-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
VQ-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
Diff-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
FontDiffuser	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
IF-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
MX-Font++	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
DA-Font	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
HF-Font (Ours)	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡
Ground Truth	蜡 劳 滥 昆	寒 好 轰 葫	或 基 加 兼	掀 局 觉 卡

793 Figure 10: Qualitative comparison results on UFSC dataset. Characters in the blue boxes suffer  
 794 from stroke errors (e.g., missing or redundant strokes), while the red boxes highlight conspicuous  
 795 blurring or artifacts. Zoom in to see details.  
 796

#### 798 A.4 MORE ABLATION STUDIES

800 Like Section 4.4, we carry out the additional ablation experiments on the UFUC dataset.

801 **The Effect of Different High-pass Filters.** We conduct ablation experiments to evaluate the im-  
 802 pacts of different high-pass filters. As illustrated in Figure 11, we compare the Laplacian filter with  
 803 three other commonly-used filters, including Wavelet, Fourier, and Sobel. From the figure, we no-  
 804 tice that the Wavelet filter could capture local high-frequency details but often yields overly sparse  
 805 and fragmented features, which may impede stroke continuity modelling. Besides, the Fourier fil-  
 806 ter tends to produce dim and blurry edges in the spatial domain. Additionally, although the Sobel  
 807 operator provides enhancement along the gradient directions, it is susceptible to noise, resulting in  
 808 disjointed contours. In contrast, the Laplacian filter extracts more accurate and sharper character  
 809 contours and structures. The quantitative results in Table 6 also confirm its advantage in extracting  
 high-frequency style features.

810  
**The Effect of GRL and ES-RPB in SCFM.** To assess the roles of Grouped Residual Layer (GRL)  
811 and Exponential-Space Relative Position Bias (ES-RPB) in our proposed Style-Content Fusion  
812 Module (SCFM), we perform a series of ablation experiments. The quantitative results are sum-  
813 marized in Table 7, and the qualitative comparisons are presented in Figure 12.  
814

815 From the results, we can find that removing GRL while retaining ES-RPB alone would lead to a  
816 performance drop, with SSIM decreasing from 0.7518 to 0.7327 and RMSE increasing from 0.2692  
817 to 0.2797. This indicates that GRL facilitates effective feature fusion via structured residual learn-  
818 ing. As shown in Figure 12, the absence of GRL results in blurred stroke details, underscoring its  
819 importance in preserving glyph topological integrity.  
820

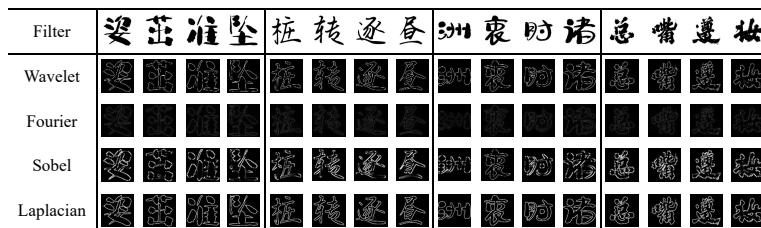
821 Similarly, disabling ES-RPB also degrades the model’s performance, notably increasing LPIPS  
822 (0.1803) and FID (56.2451), which demonstrates its efficacy in spatial alignment and geometric  
823 fidelity. Without ES-RPB, the generated glyphs tend to exhibit positional distortion and weakened  
824 visual consistency.  
825

826 **Table 6: Quantitative results on different**  
827 **high-pass filters.** It is obvious that Laplacian  
828 filter yields the best overall performance.

Filter	SSIM↑	RMSE↓	LPIPS↓	FID↓
Wavelet	0.7216	0.2987	0.1782	57.1143
Fourier	0.6479	0.3321	0.2236	66.7955
Sobel	0.6995	0.3103	0.1861	62.9451
Laplacian	<b>0.7518</b>	<b>0.2692</b>	<b>0.1544</b>	<b>50.0632</b>

829 **Table 7: Quantitative results on GRL and ES-**  
830 **RPB in SCFM.** The first row denotes the stan-  
831 **dard cross-attention mechanism.**

	Modules	SSIM↑	RMSE↓	LPIPS↓	FID↓
	GRL	ES-RPB			
×	✗	0.7152	0.3063	0.2089	60.0298
✓	✗	0.7284	0.2845	0.1803	56.2451
✗	✓	0.7327	0.2797	0.1736	54.4714
✓	✓	<b>0.7518</b>	<b>0.2692</b>	<b>0.1544</b>	<b>50.0632</b>



832 **Figure 11: Effect of different high-pass filters on style feature extraction.** Zoom in to see details.  
833

Content	跨	矿	空	俊	腊	块	客	君	肋	离	冷	泪	来	亏	拉	乐
w/o (GRL & ES-RPB)	跨	矿	空	俊	腊	块	客	君	肋	离	冷	泪	来	亏	拉	乐
w/o ES-RPB	跨	矿	空	俊	腊	块	客	君	肋	离	冷	泪	来	亏	拉	乐
w/o GRL	跨	矿	空	俊	腊	块	客	君	肋	离	冷	泪	来	亏	拉	乐
Full Model	跨	矿	空	俊	腊	块	客	君	肋	离	冷	泪	来	亏	拉	乐
Ground Truth	跨	矿	空	俊	腊	块	客	君	肋	离	冷	泪	来	亏	拉	乐

834 **Figure 12: Qualitative results on GRL and ES-RPB in SCFM.** The green boxes point details that  
835 are better generated by the full model. Zoom in to see details.  
836

## 837 838 839 840 841 842 843 A.5 THEORETICAL ANALYSES ON THE EFFECTIVENESS OF GRL AND ES-RPB IN SCFM

844 **Theoretical Explanation of GRL.** GRL includes two core components: the grouping scheme and  
845 the residual structure. The grouping scheme is proposed to efficiently reduce both parameters and  
846 computations without sacrificing representation performance. Meanwhile, the residual structure aids  
847 optimization by allowing the model to find the optimal solution in a residual space. The explanations  
848 are as follows:  
849

850 First, we analyze the reductions of parameters and Multiply-ACcumulate operations (MACs). For  
851 a conventional linear layer with both input and output dimensions equal to  $N$ , the complexity in

864 terms of parameters and MACs amounts to  $N^2$ . However, by splitting the input into two parts of  
 865 size  $\frac{N}{2}$  and processing them with two linear layers, the total complexity decreases to  $\frac{N^2}{2}$ . Hence,  
 866 the grouping scheme yields a clear halving of both the parameter count and computational cost.  
 867

868 Second, we analyze why the grouping scheme does not lead to severe performance degradation. A  
 869 possible concern about the grouping scheme is that splitting channels may reduce the interactions  
 870 between different groups of features. Here, we will show that, at least for query and key matrices  
 871 in the attention mechanism, this worry is unnecessary. Specifically, let the query and key inputs be  
 872  $X_q \in \mathbb{R}^{N_q \times C}$  and  $X_k \in \mathbb{R}^{N_k \times C}$ , where  $N_q$  and  $N_k$  denote the sequence lengths of query and key,  
 873 respectively.  $C$  is the channel dimension. Then, we divide each of them into two halves as follows:  
 874

$$X_q = [X_q^{(1)}, X_q^{(2)}] \quad X_k = [X_k^{(1)}, X_k^{(2)}] \quad (12)$$

875 where  $X_q^{(i)} \in \mathbb{R}^{N_q \times \frac{C}{2}}$  and  $X_k^{(i)} \in \mathbb{R}^{N_k \times \frac{C}{2}}$ . Then we apply grouped linear projections with parameter  
 876 matrices  $W_q^{(i)}, W_k^{(i)} \in \mathbb{R}^{\frac{C}{2} \times \frac{C}{2}}$  to derive the query matrix  $Q$  and the key matrix  $K$  as:  
 877

$$Q = [X_q^{(1)}W_q^{(1)}, X_q^{(2)}W_q^{(2)}] \quad K = [X_k^{(1)}W_k^{(1)}, X_k^{(2)}W_k^{(2)}] \quad (13)$$

880 Next, we perform the matrix product of  $Q$  and  $K^T$  by:  
 881

$$\begin{aligned} 882 \quad QK^T &= [X_q^{(1)}W_q^{(1)}, X_q^{(2)}W_q^{(2)}] \begin{bmatrix} X_k^{(1)}W_k^{(1)} \\ X_k^{(2)}W_k^{(2)} \end{bmatrix}^T \\ 883 &= X_q^{(1)}W_q^{(1)}(X_k^{(1)}W_k^{(1)})^T + X_q^{(2)}W_q^{(2)}(X_k^{(2)}W_k^{(2)})^T \end{aligned} \quad (14)$$

887 Among them, each element of  $QK^T$  is calculated as follows:  
 888

$$(QK^T)_{ij} = (X_q^{(1)}W_q^{(1)})_i \cdot (X_k^{(1)}W_k^{(1)})_j + (X_q^{(2)}W_q^{(2)})_i \cdot (X_k^{(2)}W_k^{(2)})_j \quad (15)$$

890 From Equation 15, we can see that each value is the sum of contributions from both channel groups.  
 891 Therefore, although each projection operates on a separate group of channels, the subsequent dot-  
 892 product naturally aggregates information across all groups, thereby largely preventing severe per-  
 893 formance degradation.  
 894

895 Third, we analyze the role of the residual structure in GRL. As discussed in Section 3.4, an attention  
 896 block typically involves multiple matrix multiplications, which makes it difficult for the network to  
 897 directly learn the optimal parameters. To be specific, given the input of attention mechanism as  $X$   
 898 and the desired projections as  $\hat{Q}$ ,  $\hat{K}$ , and  $\hat{V}$ , the network without residual connections must learn  
 899 these mappings entirely from scratch. In contrast, when introducing a residual structure for the QKV  
 900 linear layers, the network only needs to model  $\hat{Q} - X$ ,  $\hat{K} - X$ , and  $\hat{V} - X$ . This transformation  
 901 effectively shifts the optimization from the original linear space to a residual space, which tends  
 902 to be smoother and easier to navigate. Consequently, the residual design not only facilitates more  
 903 efficient optimization but also enhances the representational capacity of the QKV layers.  
 904

905 **Theoretical Explanation of ES-RPB.** To better understand the spatial prior introduced by ES-  
 906 RPB, we provide a visual comparison between linear and exponential coordinate mappings used in  
 907 position encoding. As illustrated in Figure 13, the exponential transformation maps the absolute  
 908 positional coordinates  $(\Delta X, \Delta Y)$  from linear space to exponential space  $(\Delta \hat{X}, \Delta \hat{Y})$ , thereby in-  
 909 creasing spatial sensitivity near the centre. This design reflects a distance-aware inductive bias that  
 910 closer pixels should contribute more to attention, which is a desirable property for spatial alignment  
 911 in the font generation task. Apart from this, ES-RPB also incorporates two other improvements.  
 912 First, it replaces direct parameter learning with a lightweight MLP to minimize the noise impact  
 913 and enhance training stability. Second, instead of learning relative positional parameters in isolation,  
 914 ES-RPB models them within the shared functional space of the MLP, which encourages richer  
 915 interactions and better generalization across spatial contexts.  
 916

## 917 A.6 MORE GENERATION RESULTS

918 To further validate the robustness and generalization capability of our proposed HF-Font, we conduct  
 919 two extra groups of experiments. The first group investigates the model’s performance across a  
 920

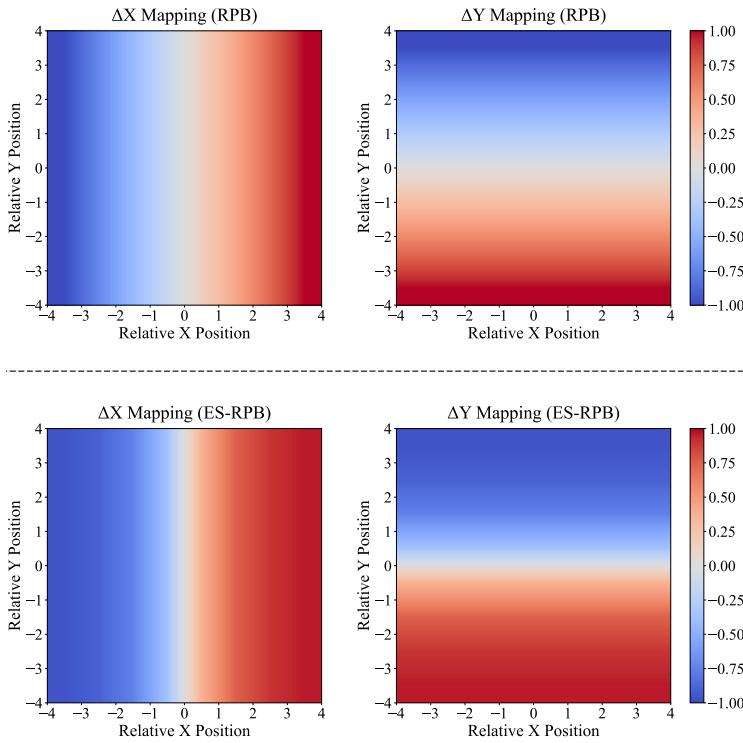


Figure 13: **Comparison between coordinate mapping functions used in RPB and ES-RPB.** The top row shows linear normalized mapping in the X-direction (left) and Y-direction (right). The bottom row displays exponential-space mapping, where changes near the centre are more prominent. The colour scale represents the relative position bias values, ranging from -1 (blue) to 1 (red).

wide range of diverse and stylistically rich typefaces, thereby aiming to assess its adaptability to substantial style variations (Figure 14). The second group targets the generation of structurally complex characters, which contain dense strokes and intricate component arrangements, to verify the model’s ability in preserving fine-grained structural fidelity under challenging cases (Figure 15). Experimental results from both settings consistently demonstrate the effectiveness and reliability of our approach in handling a broad spectrum of font generation scenarios.

#### A.7 CROSS-LINGUAL FONT GENERATION

An interesting question is whether our method can be applied to the generation of other Eastern Asian font styles, such as Japanese Kana and Korean Hangul. Obviously, collecting a large number of character images in these font styles for model training directly would fulfil this goal but also be highly time-consuming and labour-intensive. On the other hand, considering the structural and visual similarities among these font styles, we thus utilize the framework trained on our collected Chinese fonts to extend the generation of other font styles directly. To be specific, we randomly choose 15 fonts with 100 Japanese Kana characters and 10 fonts with 125 Korean Hangul characters for testing, respectively. Some generation results are displayed in Figure 16. We can see that since the structures of the Japanese Kana and Korean Hangul characters are simpler than those of Chinese characters, the generation results closely resemble the ground truths, demonstrating the robust cross-domain generalization ability of our proposed model.

#### A.8 USAGE OF LARGE LANGUAGE MODELS

During manuscript preparation, we employed several large language models (such as ChatGPT and DeepSeek) solely for the purpose of polishing the writing and improving readability. These models were not used for research ideation, experiment design, data analysis, or result interpretation. The authors take full responsibility for the content of the paper.

972  
973  
974  
975  
976  
977

Content	賴覽 兰况 款苦	立类牢扣克孔	帘良荔愣梨联	两老例连丽吏
Generated Results	賴覽 兰况 款苦	立类牢扣克孔	帘良荔愣梨联	两老例连丽吏
Ground Truth	賴覽 兰况 款苦	立类牢扣克孔	帘良荔愣梨联	两老例连丽吏
Content	坑嘵廊李了里	几回昏化卉吉	劣将坚轿桔荐	练亮疗怜粒恋
Generated Results	坑嘵廊李了里	几回昏化卉吉	劣将坚轿桔荐	练亮疗怜粒恋
Ground Truth	坑嘵廊李了里	几回昏化卉吉	劣将坚轿桔荐	练亮疗怜粒恋
Content	帝測炒碘带财	景姐见讲杰决	革丰肤父芬盖	可看开凯具刊
Generated Results	帝測炒碘带财	景姐见讲杰决	革丰肤父芬盖	可看开凯具刊
Ground Truth	帝測炒碘带财	景姐见讲杰决	革丰肤父芬盖	可看开凯具刊
Content	邻裂临料僚猎	固瓜刮怪关归	怯困宽刻考灸	副柑哥工公勾
Generated Results	邻裂临料僚猎	固瓜刮怪关归	怯困宽刻考灸	副柑哥工公勾
Ground Truth	邻裂临料僚猎	固瓜刮怪关归	怯困宽刻考灸	副柑哥工公勾
Content	汇黄宦画沪忽	奋伏负妇羔告	佳剂颊鸡祸混	吵忱持惦度饭
Generated Results	汇黄宦画沪忽	奋伏负妇羔告	佳剂颊鸡祸混	吵忱持惦度饭
Ground Truth	汇黄宦画沪忽	奋伏负妇羔告	佳剂颊鸡祸混	吵忱持惦度饭

Figure 14: **More generation results on various font styles.** The first row shows the content images, the second row denotes the characters generated by our HF-Font, and the third row represents the ground truth. Zoom in to see details.994  
995  
996  
997  
998  
999  
1000

Content	塞鲨骚蕊赛撒	燃鹊壤熔融踩	蠕褥裙趣锹儒	搂擎瘤嚷瓢擒
Generated Results	塞鲨骚蕊赛撒	燃鹊壤熔融踩	蠕褥裙趣锹儒	搂擎瘤嚷瓢擒
Ground Truth	塞鲨骚蕊赛撒	燃鹊壤熔融踩	蠕褥裙趣锹儒	搂擎瘤嚷瓢擒
Content	厦霎瞻擅裳摄	煽甥盛誓嗜匙	剩赚嘱骤撞撰	臊烧瘦淑释输
Generated Results	厦霎瞻擅裳摄	煽甥盛誓嗜匙	剩赚嘱骤撞撰	臊烧瘦淑释输
Ground Truth	厦霎瞻擅裳摄	煽甥盛誓嗜匙	剩赚嘱骤撞撰	臊烧瘦淑释输
Content	瘫檀糖藤蹄瞳	撕艘髓嗽穗酸	鼠蹋嗓肆踏嘶	蔬曙蟀霜瞬寢
Generated Results	瘫檀糖藤蹄瞳	撕艘髓嗽穗酸	鼠蹋嗓肆踏嘶	蔬曙蟀霜瞬寢
Ground Truth	瘫檀糖藤蹄瞳	撕艘髓嗽穗酸	鼠蹋嗓肆踏嘶	蔬曙蟀霜瞬寢
Content	躺腾艇褪臀豌	蔚瘟喻萎趟慰	蜀隧潭巍薇腕	睡鵝舞題雾魏
Generated Results	躺腾艇褪臀豌	蔚瘟喻萎趟慰	蜀隧潭巍薇腕	睡鵝舞題雾魏
Ground Truth	躺腾艇褪臀豌	蔚瘟喻萎趟慰	蜀隧潭巍薇腕	睡鵝舞題雾魏
Content	箫蟹薪醒熊墟	橡潇蝎懈馨歇	酗癣锈薛熏蜒	膝霞嘻嫌鑲裘
Generated Results	箫蟹薪醒熊墟	橡潇蝎懈馨歇	酗癣锈薛熏蜒	膝霞嘻嫌鑲裘
Ground Truth	箫蟹薪醒熊墟	橡潇蝎懈馨歇	酗癣锈薛熏蜒	膝霞嘻嫌鑲裘

Figure 15: **More generation results on complex characters.** The first row shows the content images, the second row denotes the characters generated by our HF-Font, and the third row represents the ground truth. Zoom in to see details.1022  
1023  
1024  
1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

Content	ぐせみむりん	ねのはほみテ	ぎざだづでス	ごじげぜぞに
Generated Results	ぐせみむりん	ねのはほみテ	ぎざだづでス	げごじぜぞに
Ground Truth	ぐせみむりん	ねのはほみテ	ぎざだづでス	げごじぜぞに
Content	めやもよらな	たちつてぬを	ヤユヨラリせ	どばびぶぼす
Generated Results	めやもよらな	たぢつてぬを	ヤユヨラリせ	どばびぶぼす
Ground Truth	めやもよらな	たちつてぬを	ヤユヨラリせ	どばびぶぼす

(a) Generated Results on Japanese (Kana) characters

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

Content	교콘콤콥꼿게	쾌팟꽝꽝꽝	갑갑갑감갑객	괴괴권권뵙가
Generated Results	교콘콤콥꼿게	쾌팟꽝꽝꽝	갑갑갑감갑객	괴괴권권뵙가
Ground Truth	교콘콤콥꼿게	쾌팟꽝꽝꽝	갑갑갑감갑객	괴괴권권뵙가
Content	구국군굳굼괌	규그극근글강	가각간간갈과	값갓갓강갓권
Generated Results	구국군굳굼괌	규그극근글강	가각간간갈과	값갓갓강갓권
Ground Truth	구국군굳굼괌	규그극근글강	가각간간갈과	값갓갓강갓권

(b) Generated Results on Korean (Hangul) characters

1061

1062

Figure 16: **Generation results on cross-lingual scenario.** In this case, we generate some Japanese Kana and Korean Hangul glyph images with the model trained on Chinese fonts.

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079