

StrokeGAN: Reducing Mode Collapse in Chinese Font Generation via Stroke Encoding

Jinshan Zeng¹, Qi Chen¹, Yunxin Liu¹, Mingwen Wang^{1*}, Yuan Yao^{2*}

¹ School of Computer and Information Engineering, Jiangxi Normal University, Nanchang, China

² Department of Mathematics, Hong Kong University of Science and Technology, Hong Kong

jinshanzeng@jxnu.edu.cn, chenqi970226@gmail.com, 201841600023@jxnu.edu.cn, mwwang@jxnu.edu.cn, yuany@ust.hk

Abstract

The generation of stylish Chinese fonts is an important problem involved in many applications. Most of existing generation methods are based on the deep generative models, particularly, the generative adversarial networks (GAN) based models. However, these deep generative models may suffer from the mode collapse issue, which significantly degrades the diversity and quality of generated results. In this paper, we introduce a one-bit stroke encoding to capture the key mode information of Chinese characters and then incorporate it into CycleGAN, a popular deep generative model for Chinese font generation. As a result we propose an efficient method called *StrokeGAN*, mainly motivated by the observation that the stroke encoding contains amount of mode information of Chinese characters. In order to reconstruct the one-bit stroke encoding of the associated generated characters, we introduce a stroke-encoding reconstruction loss imposed on the discriminator. Equipped with such one-bit stroke encoding and stroke-encoding reconstruction loss, the mode collapse issue of CycleGAN can be significantly alleviated, with an improved preservation of strokes and diversity of generated characters. The effectiveness of StrokeGAN is demonstrated by a series of generation tasks over nine datasets with different fonts. The numerical results demonstrate that StrokeGAN generally outperforms the state-of-the-art methods in terms of content and recognition accuracies, as well as certain stroke error, and also generates more realistic characters.

Introduction

The stylish Chinese font generation has attracted rising attention within recent years (Lin et al. 2016; Cha et al. 2020; Chang, Gu, and Zhang 2017; Tian 2017; Kong and Xu 2017; Jiang et al. 2017, 2019; Chang et al. 2018; Chen et al. 2019; Wu, Yang, and Hsu 2020; Gao and Wu 2020; Zhang et al. 2020), since it has a wide range of applications including but not limited to the automatic generation of artistic Chinese calligraphy (Zhao et al. 2020), art font design (Lin et al. 2014) and personalized style generation of Chinese characters (Liu, Xu, and Lin 2012).

The existing Chinese font generation methods can be generally divided into two categories. The first category is firstly to extract some explicit features such as strokes and radicals

of Chinese characters and then utilize some traditional machine learning methods to generate new characters (Xu et al. 2005; Lin et al. 2016). The quality of feature extraction plays a central role in the first category of Chinese font generation methods. However, such feature extraction procedure is usually hand-crafted, and thus time and effort consuming.

The second category of Chinese font generation methods has been recently studied in (Tian 2017; Chang et al. 2018; Chen et al. 2019; Gao and Wu 2020; Wu, Yang, and Hsu 2020; Zhang et al. 2020) with the development of deep learning (Goodfellow, Bengio, and Courville 2016), particularly the generative adversarial networks (GAN) (Goodfellow et al. 2014). Due to the powerful expressivity and approximation ability of deep neural networks, feature extraction and generation procedures can be combined into one procedure, and thus, Chinese font generation methods in the second category can be usually realized in an end-to-end training way. Instead of using the stroke or radical features of Chinese characters, the methods in the second category usually regard Chinese characters directly as images, and then translate the Chinese font generation problem into certain image style translation problem (Zhu et al. 2017; Isola et al. 2017), for which GAN and its variants are the principal techniques. However, it is well-known that GAN usually suffers from the issue of mode collapse (Goodfellow et al. 2014), that is, producing the same patterns for different inputs by generator. Such issue will significantly degrade the diversity and quality of the generated results (see, Figure 1 below). When adapted to Chinese font generation problem, the mode collapse issue will happen more frequently due to there are many Chinese characters with very similar strokes.

Due to the artificial nature of Chinese characters, the explicit stroke information contains amount of mode information of Chinese characters (see Figure 2). This is very different from natural images, which are usually regarded to be generated according to some probability distributions at some latent spaces. Inspired by this observation, in this paper, we at first introduce a one-bit stroke encoding to preserve the key mode information of a Chinese character, and then suggest certain stroke-encoding reconstruction loss to reconstruct the stroke encoding of the generated character such that the key mode information can be well preserved, and finally incorporate them into CycleGAN for Chinese font generation (Chang et al. 2018). Thus, our suggested

*Corresponding author

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Input characters 谨 畔 寥 趋 吟 艳
 CycleGAN 饽 饽 饽 饽 饽 饽
 StrokeGAN 谨 畔 寥 趋 吟 艳

Figure 1: CycleGAN for Chinese font generation suggested in (Chang et al. 2018) suffers from the mode collapse issue when translating *Black* font to *Shu* font, while the suggested StrokeGAN can effectively tackle this issue. The settings are presented in the later experiment section.

model is called *StrokeGAN*. The contributions of this paper can be summarized as follows:


- We propose an effective method called *StrokeGAN* for the generation of Chinese fonts with unpaired data. Our main idea is firstly to introduce a one-bit stroke encoding to capture the mode information of Chinese characters and then incorporate it into the training of CycleGAN (Chang et al. 2018), in the purpose of alleviating the mode collapse issue of CycleGAN and thus improving the diversity of its generated characters. In order to preserve the stroke encoding, we introduce a stroke-encoding reconstruction loss to the training of CycleGAN. By the use of such one-bit stroke encoding and the associated reconstruction loss, StrokeGAN can effectively alleviate the mode collapse issue for Chinese font generation, as shown by Figure 1.
- The effectiveness of StrokeGAN is verified over a set of Chinese character datasets with 9 different fonts (see Figure 3), that is, a handwriting font, 3 standard printing fonts and 5 pseudo-handwriting fonts. Compared to CycleGAN for Chinese font generation (Chang et al. 2018), StrokeGAN can generate Chinese characters with higher quality and better diversity, particularly, the strokes are better preserved. Besides CycleGAN (Chang et al. 2018), our method also outperforms the other state-of-the-art methods including *zi2zi* (Tian 2017) and Chinese typography transfer (CTT) method (Chang, Gu, and Zhang 2017) using paired data, in terms of generating Chinese characters with higher quality and accuracy. Some generated characters by our method with 9 different fonts can be found in Figure 3. It can be observed that these generated characters of StrokeGAN are very realistic.

Related Work

In recent years, many generation methods of stylish Chinese fonts have been suggested in the literature (Tian 2017; Chang, Gu, and Zhang 2017; Chang et al. 2018; Chen et al. 2019; Wu, Yang, and Hsu 2020; Jiang et al. 2017, 2019; Zhang et al. 2020) with the development of deep learning. In (Tian 2017), the authors adapted *pix2pix* model developed in (Isola et al. 2017) for the image style translation problem to Chinese font generation and then suggested *zi2zi* method with paired training data, that is, there is a one-to-one correspondence between the characters in the source (input) style

Stroke	Name	Example	Stroke	Name	Example
一	Heng (横)	丛	㇇	HengZhezhepie (横折折撇)	及
丨	Ti (提)	习	㇇	HengZhezheze (横折折折)	凸
丨	Shu (竖)	十	㇇	HengZhezhezhegou (横折折折钩)	乃
丨	ShuGou (竖钩)	小	㇇	HengPiewangou (横撇弯钩)	那
ノ	Pie (撇)	八	丨	ShuTi (竖提)	长
丶	Dian (点)	立	丨	ShuWan (竖弯)	四
丶	Na (捺)	人	丨	ShuWangou (竖弯钩)	儿
㇇	HengGou (横钩)	买	㇇	ShuZhe (竖折)	山
㇇	HengPie (横撇)	水	㇇	ShuZhezhe (竖折折)	鼎
㇇	HengZhe (横折)	口	㇇	ShuZhezhegou (竖折折钩)	马
㇇	HengZhegou (横折钩)	用	㇇	ShuZhepie (竖折撇)	专
㇇	HengZheti (横折提)	认	㇇	PieZhe (撇折)	云
㇇	HengZhezhe (横折折)	回	㇇	PieDian (撇点)	女
㇇	HengZhewan (横折弯)	没	㇇	WanGou (弯钩)	狗
㇇	HengZhengwangou (横折弯钩)	九	㇇	XieGou (斜钩)	我
㇇	HengXiegou (横斜钩)	飞	㇇	WuGou (卧钩)	心

(a) 32 basic strokes that make up Chinese characters



(b) Strokes of the Chinese character "Yi"

Figure 2: (a) 32 basic strokes that make up Chinese characters. The first and fourth columns present these 32 basic strokes, the second and fifth columns present the names of these basic strokes, and the third and sixth columns present some typical Chinese characters involved these basic strokes. (b) The strokes of the Chinese character "Yi".

domain and target (output) style domain. Similar idea was extended to realize the Chinese character generation from one font to multiple fonts in (Chen et al. 2019). Besides (Tian 2017) and (Chen et al. 2019), some other paired data based Chinese font generation methods were suggested in (Chang, Gu, and Zhang 2017; Jiang et al. 2017; Wu, Yang, and Hsu 2020). However, it is usually human-intensive to build up the paired training data. In order to overcome this challenge, (Chang et al. 2018) adapted CycleGAN developed in (Zhu et al. 2017) for the image style translation to Chinese font generation based on unpaired training data. Yet, the CycleGAN based method suggested in (Chang et al. 2018) (called *CCG-CycleGAN*) may suffer from the mode collapse issue (Goodfellow et al. 2014). When mode collapse occurs, the generator produces fewer patterns for different inputs, and thus significantly degrades the diversity and quality of generated results.

Motivated by the observation from traditional Chinese character generation and recognition methods (see (Xu et al. 2005; Kim, Liu, and Kim 1999)) that the explicit stroke feature can provide much mode information for a Chinese character, in this paper, we incorporate such stroke information into the training of CycleGAN for Chinese font generation (Chang et al. 2018) to tackle the issue of mode collapse, via introducing a one-bit stroke encoding and certain



Figure 3: Some Chinese characters with 9 different fonts generated by StrokeGAN. The first row presents the generated Chinese characters with the handwriting font, the second to the fourth rows present the generated Chinese characters with three different standard printing fonts, and the fifth to the ninth rows present the generated Chinese characters with five different pseudo-handwriting fonts.

stroke-encoding reconstruction loss. The very recent papers (Wu, Yang, and Hsu 2020; Jiang et al. 2019; Zhang et al. 2020) also incorporated some stroke or radical information of Chinese characters into Chinese font generation. Their main idea is firstly to utilize a deep neural network to extract the strokes or radicals of Chinese characters and then merge them by another deep neural network, which is very different to our idea of the use of a very simple one-bit stroke encoding. According to our later numerical experiments, our introduced one-bit stroke encoding is very effective.

The rest of this paper is organized as follows. In Section 2, we present some preliminary work. In Section 3, we introduce the proposed method in detail. In Section 4, we provide a series of experiments to demonstrate the effectiveness of the proposed method. We conclude this paper in Section 5.

Preliminary Work

In this section, we introduce some preliminary work, which serves as the basis of this paper.

Generative Adversarial Networks

Generative adversarial networks (GAN) (Goodfellow et al. 2014) have achieved great achievements in the task of high-quality image synthesis. The classic GAN model consists of two parts: a generator G and a discriminator D . The generator G generates fake images, and the discriminator D judges whether the images generated by G are fake or real. Mathematically, GAN can be formulated as the following two-player minimax game between G and D ,

$$\min_G \max_D \mathbb{E}_{x \sim \mathbb{P}_{data}} [\log D(x)] + \mathbb{E}_{z \sim \mathbb{P}_z} [\log(1 - D(G(z)))]$$

where \mathbb{P}_{data} and \mathbb{P}_z are the distributions of data x and the input noise variable z for the generator, respectively. In practice, the generator G and discriminator D are generally represented by some deep neural networks.

Conditional GAN

The conditional GAN (cGAN) was suggested in (Mirza and Osindero 2014) mainly to embed some conditional information such as the category information of samples. Such idea was later exploited in (Isola et al. 2017) for the image style translation. For cGAN, besides the original input z , the conditional information c is also a part of input for G . Thus, the objective for GAN should be slightly modified for cGAN, shown as follows:

$$\mathcal{L}_{adv}(D, G) = \mathbb{E}_{x \sim \mathbb{P}_{data}} [\log D(x)] + \mathbb{E}_{z \sim \mathbb{P}_z, c \sim \mathbb{P}_c} [\log(1 - D(G(z, c)))]$$

where \mathbb{P}_c represents the distribution of the referred conditional information c .

CycleGAN

The training of cGAN (Isola et al. 2017) is based on the paired data, of which the collection is usually time-consuming and laborious. In order to overcome this challenge, CycleGAN was proposed in (Zhu et al. 2017) for the image style translation based on the unpaired data. The main idea of CycleGAN is to preserve key attributes between the source and target domains by utilizing a cycle consistency loss. Specifically, let $x \in \mathcal{X}$ and $x' \in \mathcal{X}'$ be two images from two different style domains conditioned on respectively some conditional information domains $c \in \mathcal{C}$ and $c' \in \mathcal{C}'$. In order to realize the bidirectional translation between them, two generators $G : \mathcal{X} \times \mathcal{C} \rightarrow \mathcal{X}'$ and $G' : \mathcal{X}' \times \mathcal{C}' \rightarrow \mathcal{X}$ are exploited. With these, the cycle consistency loss can be defined as follows:

$$\mathcal{L}_{cyc}(G, G') = \mathbb{E}_{x \sim \mathbb{P}_x, c \sim \mathbb{P}_c, c' \sim \mathbb{P}_{c'}} [\|x - G'(G(x, c), c')\|_1] + \mathbb{E}_{x' \sim \mathbb{P}_{x'}, c' \sim \mathbb{P}_{c'}, c \sim \mathbb{P}_c} [\|x' - G(G'(x', c'), c)\|_1]$$

where $x \in \mathcal{X}$, $x' \in \mathcal{X}'$, $c \in \mathcal{C}$, $c' \in \mathcal{C}'$, and \mathbb{P}_* represents the distribution of the associated domain $*$. The generators G and G' are trained to make the cycle consistency loss $\mathcal{L}_{cyc}(G, G')$ small.

StrokeGAN for Chinese Font Generation

In this section, we describe the proposed *StrokeGAN* for automatic generation of stylish Chinese font. The core idea of StrokeGAN is to incorporate some one-bit stroke encodings of Chinese characters into CycleGAN to alleviate the issue of mode collapse, as presented in Figure 4, mainly motivated by the basic observation that the stroke information embodies amount of mode information of Chinese characters. Recent studies suggested that mode collapse in GANs might be a by-product of removing sparse outlying modes toward robust estimation (Gao et al. 2019; Gao, Yao, and Zhu 2020). So a natural strategy to alleviate mode collapse is to enforce faithful conditional distribution reconstruction, conditioning on important modes represented by stroke encoding here.

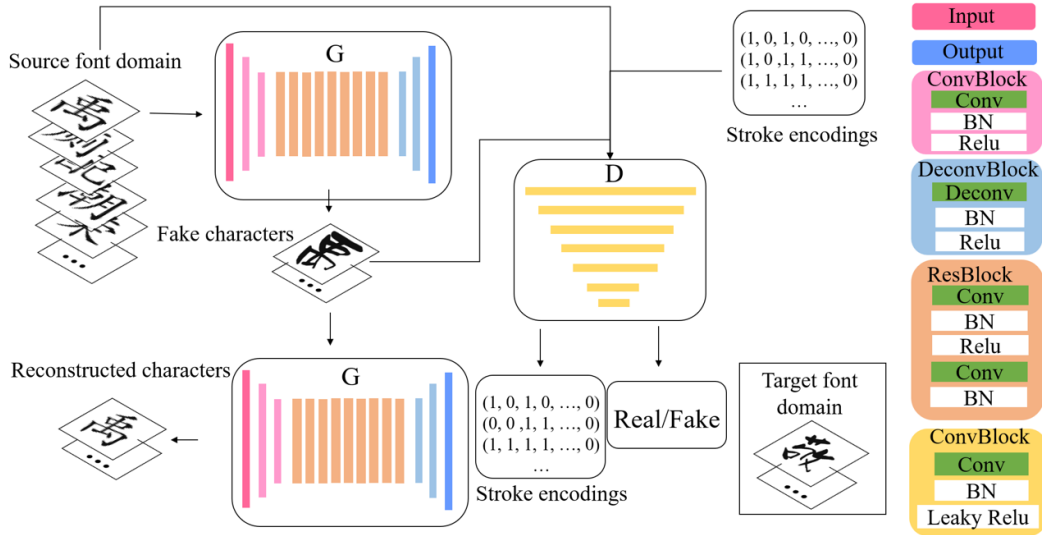


Figure 4: The diagram of StrokeGAN. There are a discriminator D and a generator G in StrokeGAN. The network architectures of generator and discriminator are stacked by modules listed in the right column.

From Figure 4, we at first yield the stroke encodings of Chinese characters via certain one-bit encoding way, and then take the characters in the source font domain as the inputs of generator G . After G , we yield a fake character in the target domain, then send such fake character to both generator and discriminator, where generator tries to reconstruct the character in the source font domain, and discriminator attempts to distinguish whether the generated character is real or fake and also reconstruct its stroke encoding. Thus, distinguished with the original CycleGAN for Chinese font generation in (Chang et al. 2018), there are two parts in the discriminator D , i.e.,

$$D : x \rightarrow (D_{src}(x), D_{st}(x)),$$

where $D_{src}(x)$ and $D_{st}(x)$ represent respectively the probability distribution over the source font domain and its stroke encoding for a given character x .

Stroke Encoding

From Figure 4, the stroke encoding of character is taken as a part of input of CycleGAN. To realize this, we introduce a simple one-bit encoding way to yield the associated stroke encoding for a given Chinese character. Specifically, according to Figure 2, there are in total 32 kinds of strokes to make up Chinese characters. Thus, for any given Chinese character x , we define its stroke encoding as a 32-dimensional vector $c \in \{0, 1\}^{32}$ with the i -th entry being 1 if the i -th kind of stroke is included in x and otherwise 0 for all i from 1 to 32. In this paper, we only use the indicator function of such kind of stroke instead of its exact number for a given Chinese character mainly in consideration of the robustness of StrokeGAN. This is in general sufficient according to our later experiments (see, Figure 7).

Training Loss for StrokGAN

The training loss for StrokeGAN consists of three parts, that is, the general adversarial loss, cycle consistency loss and stroke-encoding reconstruction loss, where the *stroke-encoding reconstruction loss* is firstly introduced in this paper for the generation of stylish Chinese fonts.

The first part of loss is the adversarial loss defined commonly as follows,

$$\mathcal{L}_{adv}(D, G) = \mathbb{E}_x [\log D_{src}(x)] + \mathbb{E}_x [\log(1 - D_{src}(G(x)))] \quad (1)$$

where the generator G generates the fake character $G(x)$ conditional over the input character x , and the first part of discriminator D_{src} attempts to distinguish the generated character is real or fake.

The second part of loss is the cycle consistency loss, which is introduced to let the generator reconstruct the character in the source font domain from the generated fake character, and thus avoid using the paired data. Specifically, such part of loss can be defined as follows:

$$\mathcal{L}_{cyc}(G) = \mathbb{E}_x [\|x - G(G(x))\|_1], \quad (2)$$

where $G(x)$ represents the generated fake character according to the input x , and $G(G(x))$ represents the reconstructed character from the character $G(x)$ generated by G .

Notice that in both adversarial loss and cycle consistency loss, the characters are regarded as images during the training, while the stroke information is not paid much attention. Actually, as discussed before, the stroke information embodies amount of mode information of a Chinese character, and thus is some kind of very important information that should be preserved during the training. Also, the stroke information is very special for Chinese characters and makes the generation of Chinese characters very different from the

Character style translation	Content accuracy (%) \uparrow		Recognition accuracy (%) \uparrow		Stroke error ($\times 10^{-2}$) \downarrow	
	CycleGAN	StrokeGAN	CycleGAN	StrokeGAN	CycleGAN	StrokeGAN
<i>Regular Script</i> \rightarrow <i>Shu</i>	89.56	90.48	89.36	90.52	6.79	5.69
<i>Regular Script</i> \rightarrow <i>Huawen Amber</i>	86.88	88.68	87.56	88.92	8.71	7.20
<i>Regular Script</i> \rightarrow <i>Hanyi Lingbo</i>	87.64	88.24	87.48	88.32	7.90	6.81
<i>Regular Script</i> \rightarrow <i>Imitated Song</i>	90.28	91.72	90.84	91.60	7.33	5.55
<i>Regular Script</i> \rightarrow <i>Handwriting</i>	87.08	87.64	87.12	87.60	7.67	6.50
<i>Hanyi Thin Round</i> \rightarrow <i>Hanyi Doll</i>	87.00	87.60	86.92	87.80	7.69	6.52
<i>Black</i> \rightarrow <i>Hanyi Thin Round</i>	86.60	87.76	86.60	87.92	7.79	6.96
<i>Imitated Song</i> \rightarrow <i>Black</i>	89.56	88.96	89.44	89.12	8.06	6.72
<i>Imitated Song</i> \rightarrow <i>Regular Script</i>	89.96	90.32	90.04	90.28	8.49	6.28

Table 1: The performance of StrokeGAN in 9 generation tasks via comparing with CycleGAN (Chang et al. 2018).

style translation of natural images. Specifically, the *stroke-encoding reconstruction loss* can be defined as follows:

$$\mathcal{L}_{st}(D) = \mathbb{E}_{x,c} [\|D_{st}(G(x)) - c\|_2]. \quad (3)$$

where $D_{st}(G(x))$ is the stroke encoding yielded by discriminator for the generated fake character $G(x)$. Such stroke-encoding reconstruction loss is used to guide the networks to reconstruct the stroke encodings as accurately as possible so that the modes of characters can be preserved much better.

Combining the above three parts of loss, i.e., (1)-(3), the total training loss of StrokeGAN is presented as follows:

$$\mathcal{L}_{strokegan}(D, G) = \mathcal{L}_{adv}(D, G) + \lambda_{cyc}\mathcal{L}_{cyc}(G) + \lambda_{st}\mathcal{L}_{st}(D), \quad (4)$$

where λ_{cyc} and λ_{st} are two penalty parameters. Based on the above defined loss $\mathcal{L}_{strokegan}(D, G)$, the discriminator D attempts to maximize it while the generator G tries to minimize it, shown as follows:

$$\min_G \max_D \mathcal{L}_{strokegan}(D, G). \quad (5)$$

Numerical Experiments

In this section, we provide a series of experiments to demonstrate the effectiveness of the suggested StrokeGAN. All experiments were carried out in Pytorch environment running Linux, AMD(R) Ryzen 7 2700x eight-core processor $\times 16$ CPU, GeForce RTX 2080 GPU. Our codes are available in <https://github.com/JinshanZeng/StrokeGAN>.

Experiment Settings

Datasets. The dataset used in this paper consists of 9 sub-datasets with different fonts divided into three categories, i.e., a handwriting font, three standard printing fonts $\{\textit{Black}, \textit{Regular Script}, \textit{Imitated Song}\}$, and five pseudo-handwriting fonts $\{\textit{Huawen Amber}, \textit{Shu}, \textit{Hanyi Lingbo}, \textit{Hanyi Doll}, \textit{Hanyi Thin Round}\}$, where the pseudo-handwriting fonts are personalized fonts designed by artistic font designers.

The first kind of dataset related to the handwriting Chinese characters is built up from *CASIA-HWDB1.1*¹, which

¹<http://www.nlpr.ia.ac.cn/databases/handwriting/Home.html>

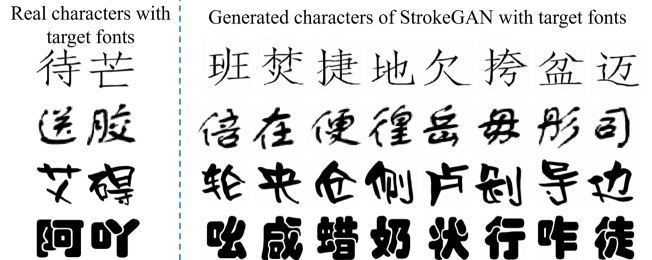


Figure 5: Examples of Chinese characters generated by StrokeGAN. The characters in the left-hand side are real characters in the target font domains, while the characters in the right-hand side are generated characters in the associated target font domains. From the first row to the fourth row, the target fonts are $\{\textit{Imitated Song}, \textit{Shu}, \textit{Hanyi Lingbo}, \textit{Huawen Amber}\}$, respectively.

was collected by 300 people. There are in total 3755 different commonly used Chinese characters written by everyone, and thus there are 3755×300 handwriting characters in this dataset. To build up our handwriting dataset, for each character, we randomly selected one sample from these 300 samples. Therefore, the size of the handwriting font dataset used in this paper is 3755. Except the handwriting Chinese characters, the other font datasets were collected by ourselves from the internet² and made automatically via TTF. Specifically, the second kind of datasets contain three printing font datasets, of which the sizes are 2560, 3757 and 2506 respectively for the *Black*, *Regular Script* and *Imitated Song* fonts. The third kind of datasets consist of five pseudo-handwriting fonts, of which the sizes are 3596, 2595, 3673, 3213 and 2840 respectively for the *Huawen Amber*, *Shu*, *Hanyi Lingbo*, *Hanyi Doll* and *Hanyi Thin Round* fonts. The size of each character was resized to $128 \times 128 \times 3$. In our experiments, we used 90% and 10% of the samples respectively as the training and test sets.

Network Structure. The network structure of the generator in StrokeGAN is the same as CycleGAN (Zhu et al. 2017), including 2 convolutional layers in the down-

²say, <http://fonts.mobanwang.com/>

CycleGAN StrokeGAN CycleGAN StrokeGAN CycleGAN StrokeGAN




Figure 6: Stoke-missing phenomena in the generated characters of CycleGAN (Chang et al. 2018). The suggested StrokeGAN can preserve strokes much better. From left to right respectively present the generated characters from *Regular Script* to *Shu*, from *Regular Script* to *Huawen Amber*, and from *Hanyi Thin Round* to *Hanyi Doll*.



Figure 7: Generated results of StrokeGAN for some very similar Chinese characters with the same stroke encodings.

sampling module, 9 residual modules with 2 convolutional layers of residual networks for each residual module and 2 deconvolutional layers in the up-sampling module, as presented in Appendix. The network structure of the discriminator in StrokeGAN is similar to PatchGAN (Isola et al. 2017) with 6 hidden convolutional layers and 2 convolutional layers in the output module, as presented in Appendix. Moreover, the batch normalization (Ioffe and Szegedy 2015) was used in all layers.

In our experiments, we used the popular Adam algorithm (Kingma and Ba 2014) as the optimizer with the associated parameters (0.5, 0.999) in both the generator and discriminator optimization subproblems. The penalty parameters of the cycle consistency loss and stroke reconstruction loss were fine-tuned at 10 and 0.18, respectively.

Evaluation Metrics. To evaluate the performance of StrokeGAN, we introduced three evaluation metrics. The first one is the commonly used *content accuracy* (Zhu et al. 2017), which was suggested to justify the quality of the contents of generated characters. Specifically, a pre-trained HCCG-GoogLeNet (Szegedy et al. 2014) was exploited to calculate the content accuracy. Besides the *content accuracy*, we also suggested the *recognition accuracy* and *stroke error* particularly for the Chinese character generation task. The *recognition accuracy* is defined as the ratio of those gen-

erated characters that can be correctly recognized by people to all the generated characters for testing, via a crowdsourcing way. Specifically, we randomly invited five Chinese adults to recognize the generated Chinese characters, and then took their average as the *recognition accuracy*. The *stroke error* is defined as the ratio of the number of missing and redundant strokes in the generated Chinese characters to its true total number of strokes. Thus, the smaller stroke error means the strokes are preserved better.

Experiment Results

Our experiments consist of three parts, where the first two parts were conducted to respectively show the effectiveness and sufficiency of the introduced one-bit stroke encoding, and the third part was conducted to demonstrate the effectiveness of StrokeGAN via comparing with the state-of-the-art methods including CycleGAN (Chang et al. 2018), zi2zi (Tian 2017) and Chinese typography transfer (CTT) method (Chang, Gu, and Zhang 2017), where the latter two methods are based on paired training data.

Effectiveness of One-bit Stroke Encoding. In these experiments, we verified the feasibility and effectiveness of our idea that incorporating the stroke information into CycleGAN can preserve the modes of Chinese characters better and thus alleviate the issue of mode collapse. In order to verify this, we implemented nine one-to-one Chinese character style translation experiments as listed in Table 1. For each experiment, we selected one font domain (say *Regular Script*) from these nine different font styles as the source font domain and another font domain (say *Shu*) as the target font domain, and then trained StrokeGAN as well as CycleGAN as the baseline. The performance of the suggested StrokeGAN is presented in Table 1. Some examples of the generated Chinese characters by StrokeGAN are shown in the previous Figure 3 and Figure 5, which demonstrate that StrokeGAN can generate very realistic Chinese characters for all these nine fonts.

From Table 1, StrokeGAN outperforms CycleGAN (without using the stroke encoding) in most of the tasks except the style translation task from *Imitated Song* to *Regular Script*. In terms of the stroke error presented in the last two columns, StrokeGAN consistently improves the performance of CycleGAN (Chang et al. 2018). This shows the feasibility and effectiveness of the suggested StrokeGAN.

Moreover, as demonstrated in the previous Figure 1, CycleGAN sometimes suffers from the issue of mode collapse (particularly, when applied to the generation task from *Black* font to *Shu* font), while the suggested StrokeGAN can significantly alleviate this issue. Besides the issue of mode collapse, as shown in Figure 6, CycleGAN also sometimes suffers from another issue, that is, missing some key strokes of Chinese characters generated, which may result in the recognition issues of these characters, while the suggested StrokeGAN can preserve these strokes much better.

Sufficiency of One-bit Stroke Encoding. Notice that there are many very similar Chinese characters having the same stroke encodings. In order to show the sufficiency of the introduced one-bit encoding, we tested the performance

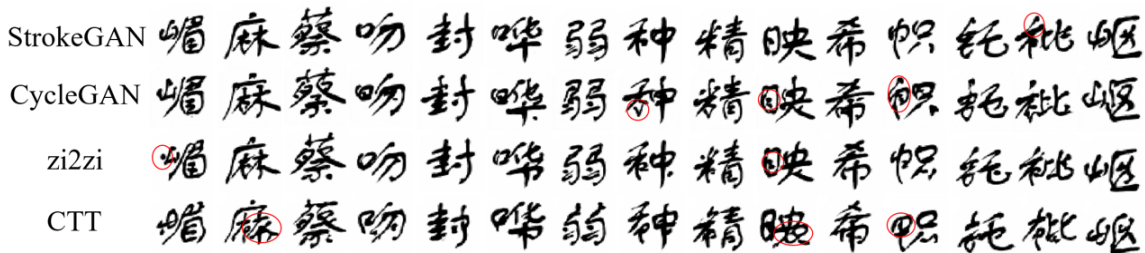


Figure 8: Some examples of Chinese characters generated by StrokeGAN and three existing methods in the generation task from *Regular Script* font to *Shu* font.

Methods	StrokeGAN	CycleGAN	zi2zi	CTT
Content acc. (\uparrow)	90.48	89.56	90.12	87.92
Recog. acc. (\uparrow)	90.52	89.36	90.04	88.01
Stroke error (\downarrow)	5.67	6.79	6.36	7.52

Table 2: Comparison on the performance of the suggested StrokeGAN with existing methods in the generation task from *Regular Script* font to *Shu* font. The first, second and third rows respectively present the content accuracies (%), recognition accuracies (%) and stroke errors in the scale $\times 10^{-2}$ of all these four methods.

of StrokeGAN on some very similar Chinese characters, as shown in Figure 7, which shows that the generated characters can be well distinguished with well preserved strokes. This demonstrates that such one-bit encoding way is in general enough to preserve the key modes of Chinese characters.

Comparison with the State-of-the-art Methods. In this part of experiments, besides CycleGAN (Chang et al. 2018), we compared StrokeGAN with two paired data based methods, i.e., *zi2zi* (Tian 2017) and CTT method (Chang, Gu, and Zhang 2017). In order to implement both *zi2zi* and CTT methods, we manually built up two paired datasets based on the *Regular Script* font and *Shu* font datasets. Since the collection of paired training data is very costly, in these experiments, we only considered the generation task from the *Regular Script* font to *Shu* font for all these four methods, while for the other generation tasks, it can be implemented similarly. The performance of these four methods is presented in Table 2, and some examples of generated characters are shown in Figure 8. From Table 2, the suggested StrokeGAN outperforms all these existing methods in terms of the suggested three evaluation metrics, and also generates the Chinese characters with the highest quality among all these methods. These experiment results demonstrate the effectiveness of the proposed StrokeGAN.

Conclusion

This paper proposes an effective Chinese font generation method called *StrokeGAN* by incorporating a one-bit stroke encoding into CycleGAN to tackle the mode collapse issue. The key intuition of our idea is that the stroke encodings of Chinese characters contain amount of mode information of

Chinese characters, unlike the natural images. A new stroke-encoding reconstruction loss was introduced to enforce a faithful reconstruction of the stroke encoding as accurately as possible and thus preserve the mode information of Chinese characters. Besides the commonly used content accuracy, the crowdsourcing recognition accuracy and stroke error are also introduced to evaluate the performance of our method. The effectiveness of StrokeGAN is demonstrated by a series of Chinese font generation tasks over 9 datasets with different fonts, comparing with CycleGAN and other two existing methods based on the paired data. The experiment results show that StrokeGAN helps preserve the stroke modes of Chinese characters in a better way and generates very realistic characters with higher quality. Besides Chinese font generation, our idea of the one-bit stroke encoding can be easily adapted to other deep generative models and applied to the font generation related to other languages such as Korean and Japanese.

Acknowledgments

The work is supported by NNSF of China (Nos. 61977038, 61876074), Thousand Talents Plan of Jiangxi Province (No. jxsq2019201124), HKRGC 16303817, and ITF UIM/390, as well as awards from Tencent AI Lab and Si Family Foundation.

References

- Cha, J.; Chun, S.; Lee, G.; Lee, B.; Kim, S.; and Lee, H. 2020. Few-shot Compositional Font Generation with Dual Memory. In *Proc. the 16th European Conference on Computer Vision (ECCV)*.
- Chang, B.; Zhang, Q.; Pan, S.; and Meng, L. 2018. Generating Handwritten Chinese Characters using CycleGAN. In *Proc. Winter Conference on Applications of Computer Vision (WACV'18)*, 1–9. Lake Tahoe, USA.
- Chang, J.; Gu, Y.; and Zhang, Y. 2017. Chinese Typography Transfer. *arXiv:1707.04904* 1–7.
- Chen, J.; Ji, Y.; Chen, H.; and Xu, X. 2019. Learning One-to-Many Stylised Chinese Character Transformation and Generation by Generative Adversarial Networks. *IET Image Processing* 13: 2680–2686.
- Gao, C.; Liu, J.; Yao, Y.; and Zhu, W. 2019. Robust estimation and generative adversarial nets. In *International Con-*

- ference on Learning Representations (ICLR), New Orleans, Louisiana, United States. ArXiv preprint arXiv:1810.02030.
- Gao, C.; Yao, Y.; and Zhu, W. 2020. Generative Adversarial Nets for Robust Scatter Estimation: A Proper Scoring Rule Perspective. *Journal of Machine Learning Research* 21(160): 1–48. ArXiv preprint arXiv:1903.01944.
- Gao, Y.; and Wu, J. 2020. GAN-Based Unpaired Chinese Character Image Translation via Skeleton Transformation and Stroke Rendering. In *Proc. the 34th AAAI Conference on Artificial Intelligence*, 646–653. New York, USA.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MA: MIT Press.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; and Xu, B. 2014. Generative Adversarial Nets. In *Proc. the 28th Conference on Neural Information Processing Systems (NeurIPS'14)*, 1–9. Montreal, Canada.
- Ioffe; and Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv:1502.03167* 1–11.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *Proc. the 30th IEEE Conference on Computer Vision and Pattern Recognition (CVPR'17)*, 1–17. Honolulu, Hawaii.
- Jiang, Y.; Lian, Z.; Tang, Y.; and Xiao, J. 2017. DCFont: An End-to-End Deep Chinese Font Generation System. In *Proc. the Association for Computing Machinery*, 1–4. New York, USA.
- Jiang, Y.; Lian, Z.; Tang, Y.; and Xiao, J. 2019. SC-Font: Structure-Guided Chinese Font Generation via Deep Stacked Networks. In *Proc. the 35rd AAAI Conference on Artificial Intelligence*, 4015–4022. Hawaii, USA.
- Kim, I.-J.; Liu, C.-L.; and Kim, J.-H. 1999. Stroke-guided pixel matching for handwritten Chinese character recognition. In *Proceed. of the Fifth International Conference on Document Analysis and Recognition (ICDAR'99)*. Bangalore, India.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. In *Proc. International Conference for Learning Representations (ICLR'14)*, 1–15. Banff, Canada.
- Kong, W.; and Xu, B. 2017. Handwritten Chinese character generation via conditional neural generative models. In *Proc. the 31st Conference on Neural Information Processing Systems (NeurIPS'17)*, 1–7. Long Beach, CA.
- Lin, J.-W.; Hong, C.-Y.; Chang, R.-I.; and Wang, Y.-C. 2016. Complete Font Generation of Chinese Characters in Personal Handwriting Style. In *Proc. the International Conference on Computing and Communications Conference*, 1–5. Nanjing, China.
- Lin, J.-W.; Wang, C.-Y.; Ting, C.-L.; and Chang, R.-I. 2014. Font Generation of Personal Handwritten Chinese Characters. In *Proc. International Conference on Graphic and Image Processing (ICGIP'14)*, 1–6. Beijing, China.
- Liu, P.; Xu, S.; and Lin, S. 2012. Automatic generation of personalized Chinese handwriting characters. In *Proc. the 4th International Conference on Digital Home (ICDH'12)*, 421–425. Guangzhou, China.
- Mirza, M.; and Osindero, S. 2014. Conditional generative adversarial nets. *arXiv:1411.1784* 1–7.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2014. Going Deeper with Convolutions. In *Proc. the 27th IEEE Conference on Computer Vision and Pattern Recognition (CVPR'14)*, 1–9. Columbus, Ohio.
- Tian, Y. 2017. Zi2zi: Master Chinese Calligraphy with Conditional Adversarial Networks. <https://github.com/kaonashityc/zi2zi>.
- Wu, S.-J.; Yang, C.-Y.; and Hsu, J. Y. 2020. CalliGAN: Style and Structure-aware Chinese Calligraphy Character Generator. *ArXiv:2005.12500* 1–8.
- Xu, S.; Lau, F. C.; Cheung, K.-W.; and Pan, Y. 2005. Automatic generation of artistic Chinese calligraphy. *IEEE Intelligent Systems* 20(3): 32–39.
- Zhang, J.; Chen, D.; Han, G.; Li, G.; He, J.; Liu, Z.; and Ruan, Z. 2020. SSNet: Structure-Semantic Net for Chinese typography generation based on image translation. *Neurocomputing* 371: 15–26.
- Zhao, B.; Tao, J.; Yang, M.; Tian, Z.; Fan, C.; and Bai, Y. 2020. Deep Imitator: Handwriting Calligraphy Imitation via Deep Attention Networks. *Pattern Recognition* 104: 1–14.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *Proc. the 30th IEEE Conference on Computer Vision and Pattern Recognition (CVPR'17)*, 2223–2232. Honolulu, Hawaii.