

HWD: A Novel Evaluation Score for Styled Handwritten Text Generation

Vittorio Pippi
vittorio.pippi@unimore.it

Fabio Quattrini
fabio.quattrini@unimore.it

Silvia Cascianelli
silvia.cascianelli@unimore.it

Rita Cucchiara
rita.cucchiara@unimore.it

University of Modena and Reggio
Emilia
Modena, IT

Abstract

Styled Handwritten Text Generation (Styled HTG) is an important task in document analysis, aiming to generate text images with the handwriting of given reference images. In recent years, there has been significant progress in the development of deep learning models for tackling this task. Being able to measure the performance of HTG models via a meaningful and representative criterion is key for fostering the development of this research topic. However, despite the current adoption of scores for natural image generation evaluation, assessing the quality of generated handwriting remains challenging. In light of this, we devise the Handwriting Distance (HWD), tailored for HTG evaluation. In particular, it works in the feature space of a network specifically trained to extract handwriting style features from the variable-length input images and exploits a perceptual distance to compare the subtle geometric features of handwriting. Through extensive experimental evaluation on different word-level and line-level datasets of handwritten text images, we demonstrate the suitability of the proposed HWD as a score for Styled HTG. The pretrained model used as backbone will be released to ease the adoption of the score, aiming to provide a valuable tool for evaluating HTG models and thus contributing to advancing this important research area.

1 Introduction

Styled Handwritten Text Generation (Styled HTG) entails producing realistic images of arbitrary handwritten text in a desired style given in the form of one or more exemplar style images. Those images can be used to: train models for document analysis tasks (*e.g.*, Handwritten Text Recognition [4, 5, 6, 9, 10, 12, 22, 24, 35, 42]) in low-resource scenarios such as ancient languages or documents by specific authors; enhance the user experience in augmented reality scenarios and the public engagement at GLAM institutions (galleries, libraries, archives, and museums); assist physically impaired people in taking notes on electronic devices. Although HTG is receiving increased interest in recent years [8, 16, 23, 34], a precise standard

for evaluation has not yet been defined. It is important to note that evaluating the similarity between two writers' calligraphy involves more than considering the overall appearance of the text images. This concerns color and texture of background and ink, stroke thickness, slant, and roundness. Nonetheless, handwriting is characterized by the shape of individual characters, ligatures, and the spacing between characters in the text. Hence, a thorough evaluation procedure should consider all these factors to ensure an accurate and meaningful assessment of the ability of Styled HTG models to imitate a desired handwriting.

The commonly adopted approach is the one proposed by Kang et al. [23], which entails exploiting the Fréchet Inception Distance (FID) [24]. The FID is computed on the features extracted by an Inception-v3 ConvNet trained on natural images from ImageNet. Thus, the FID is somehow unsatisfactory for measuring the faithfulness in the handwriting style of the generated images but rather captures the overall appearance [23, 24, 34]. Another critical point of applying the FID in the text images domain is that it leverages a backbone trained on images whose aspect ratio is very different from text images. These latter are usually wider than high, while natural images in ImageNet are roughly squared. For this reason, in HTG evaluation, the FID is commonly computed on the beginning part of the text image, discarding the rest. This approach offers invariance with respect to the textual content but is prone to miss artifacts and dissimilarities that appear in the center or at the rightmost part of the image. Therefore, while the FID can help evaluate certain aspects of HTG, it does not provide a complete picture of the quality of the generated handwriting images. Moreover, different studies [4, 11] have shown that the FID is biased on the number of samples used for its calculation, resulting in lower values with more instances and higher and more unstable values when reducing the number of examples.

To address these challenges, we propose a new evaluation score called Handwriting Distance (HWD). The main characteristics of HWD consist of: (1) the use of robust style features extracted by a backbone trained on a large dataset of synthetic text images; (2) the application of a perceptual distance for style comparison; (3) the ability to handle variable-length text images; (4) the numerical stability even when computed on a limited number of samples. This is particularly useful in Styled HTG, where only a few real images per author are usually available. To assess the suitability of the proposed score for the Styled HTG task, we examine the values obtained when comparing sets of text images in the same style with respect to those obtained when comparing images in different styles. We demonstrate that the HWD effectively captures differences in the handwriting and is numerically stable. This makes it more suitable than the FID in expressing the performance of the Styled HTG models. Overall, the HWD could contribute to the field of HTG by providing a tailored and practical evaluation score to measure the realism and faithfulness of generated text images. The code and the weights of the convolutional backbone used to compute the HWD score can be found here: <https://github.com/aimagelab/HWD>.

2 Related Work

The early-proposed approaches to HTG [20, 39] apply handcrafted geometric statistical-based feature extraction on human-made glyphs segmentations, then combine them with appropriate ligatures and render the results with texture and background blending. The two major limitations of these approaches are their inability to render glyphs and ligatures not observed for each style and their reliance on costly human intervention. In contrast, recent deep learning-based HTG approaches can infer styled glyphs even when they are not ex-

explicitly shown in the style examples. A majority of learning-based solutions are based on GANs [19], either unconditioned (for Non-Styled HTG) or conditioned on a variety of handwriting style examples (Styled HTG). In the second scenario, style samples may consist of whole sentences or lines [13], a few words [9, 23, 34], or a single word [17, 18, 27].

The first learning-based Non-Styled HTG approach, proposed in [10], entails generating fixed-sized images conditioned on the embedding of the desired textual content but does not control the calligraphic style of the output. Since, different from natural images, handwritten text images are highly variable-sized, subsequent approaches [16] entail concatenating character images. Styled HTG approaches condition the generation on both the text content and a vector representation of the style [13, 17, 18, 23, 24, 31]. The two representations are obtained separately and then combined for generation, preventing those approaches from effectively capturing local writing style and patterns. On the other hand, the Transformer-based [33] approach adopted in [9, 34] exploits the cross-attention mechanism between the style vector representation and the content text representation to entangle the content and the style, thus better rendering local style patterns.

HTG Evaluation. As for the performance evaluation, models for HTG are evaluated by exploiting the FID [21]. The FID is a commonly used score for evaluating the quality of generative models. It exploits image representations extracted from an Inception-v3, which are fit to two multivariate Gaussian distributions, one from real images and the other from generated images. For this reason, the FID tends to focus more on general image characteristics rather than the shape of handwriting. Furthermore, the backbone network used to compute the FID is trained on ImageNet, which contains natural images whose domain and aspect ratio are completely different from those of handwritten text images, which can result in misleading values. Other adopted metrics are the Geometric Score [25] and the Character Error Rate. The latter measures the readability of the generated text images, which serves as a proxy to express their realism, *i.e.*, how similar to a well-formed text they look. However, as discussed in [23, 24, 34], these measures fail to capture all the desired characteristics of a well-generated styled text image. In this work, we propose a score specific for evaluating HTG models to address this point.

3 Handwriting Distance

Inspired by the strategy adopted for natural image generation evaluation, Styled HTG works employ the FID score, adapted as described in [23] and depicted in Figure 1 (bottom). In this work, we devise an alternative score for evaluating the performance of Styled HTG models, called HWD. The main characteristics of HWD are the domain-aware image representation strategy and the use of a perceptual distance instead of a distribution distance. The pipeline for computing our proposed score is described below and depicted in Figure 1 (top).

3.1 Text Images Representation

When evaluating Styled HTG, we consider the set of real images $\mathbf{X}_m = \{\mathbf{x}_{m,i}\}_{i=0}^N$, where N is the number of samples for writer m , and the set $\mathbf{X}'_m = \{\mathbf{x}'_{m,i}\}_{i=0}^{N'}$ of generated images in the style of writer m . In general, the number of generated images, N' , can differ from N .

Domain-Specific Feature Extraction. Given the constrained domain of the images for the HTG task, we propose to use a domain-specific backbone as feature extractor. In particular, we adopt a VGG16 pretrained on Font², a large synthetic dataset of text images built

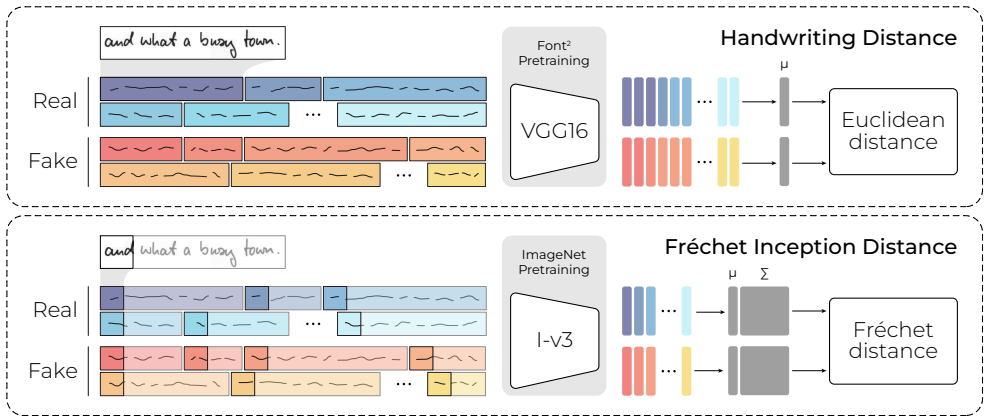


Figure 1: Pipelines for the computation of the proposed HWD score (top) and the commonly used FID score (bottom) from a set of real and generated text images with varying widths. Top: the samples are processed by a VGG16 pretrained on a synthetic dataset of text images. The different widths of the inputs result in a different number of feature vectors. The value of HWD is the Euclidean distance between the feature vectors’ means μ . Bottom: the text images are cropped to a square and fed to an Inception-v3 pretrained on ImageNet, obtaining one feature vector per image. The FID is the Fréchet distance between the distributions of the features, represented as multivariate Gaussians with mean μ and covariance matrix Σ .

according to the procedure presented in [5]. We choose VGG16 as a backbone for its superiority in extracting meaningful style representations compared to deeper networks featuring skip connections [40]. The pretraining dataset contains more than 100M samples obtained by rendering 10400 English words in 10400 calligraphic fonts and superimposing them to paper-like backgrounds. Random geometric distortions and color transformations are also performed to increase the realism of the images. We train the VGG16 backbone to classify the images according to their calligraphic font. The high visual variability of the datasets forces the network to learn features that represent the handwriting given by the font and disregard the overall visual appearance. In this way, the features extracted by our adopted backbone are strong representations of the handwriting style. For computing the HWD, we feed the pretrained VGG16 with the whole image resized to a height of 32 pixels while keeping the aspect ratio. Note that since the images have different widths, we feed them to the network one by one to avoid using padding.

Variable-length Images Representation. To handle the difference in the aspect ratio of the natural images in ImageNet and text images, the strategy adopted in [23] and the following HTG approaches entails feeding to the Inception-v3 backbone only a squared, truncated sample. The text image is reshaped to a height of 32 pixels, keeping the aspect ratio, and then cropped at the beginning to have a fixed width (32 pixels). This strategy has two drawbacks. First, it disregards a large part of the text image. Second, the characters appearing more often at the beginning of words due to language regularity assume more importance than others in computing the score. Consequently, this strategy prevents evaluating the ability of the HTG models to generate images of variable-length texts with the same quality throughout the entire image. To overcome these limitations, we propose processing the whole images with the VGG16 backbone for the HWD. The representations are obtained from the feature maps of the last convolutional block. Such feature maps have shape $1 \times W \times 512$, where W depends

on the input text image width. From those maps, we obtain a set of 512-sized feature vectors that represent the images. As a result, wide images are represented by a larger set of vectors than short ones. When computing the HWD for writer m , all the vectors from all the real images sets are gathered and then averaged (the same applies to the generated images). In the gathering set, there will be more vectors coming from the longer words, and thus, the longer words will have a bigger impact on the mean of the gathering set. In summary, each image for writer m , both real and generated, is represented by a set of vectors. The real images are represented by sets as $\mathbf{y}_{m,i} = \{f(\mathbf{x}_{m,i})_j\}_{j=0}^{W_i}$, and the generated ones by $\mathbf{y}'_{m,i} = \{f(\mathbf{x}'_{m,i})_j\}_{j=0}^{W'_i}$, where $f(\cdot)$ denotes the feature extraction, and W_i and W'_i are the number of vectors extracted from the i -th real and generated images, respectively.

3.2 Perception-Aware Feature Distance

The main idea behind the distribution distance-based evaluation scores, such as the FID, is to evaluate the performance of a generative model by its ability to generate images that match the distribution of the real ones. Our domain of interest is more constrained than the generation of natural images. In fact, Styled HTG entails considering the handwriting, expressed by subtle geometric features other than macroscopic texture. In light of this, we argue that a score capturing the perceptual aspects is more suitable than one based on the distance between feature distributions. Therefore, we employ the Euclidean distance between the averaged feature vectors of the real and generated images in the style of the same writer:

$$Y_m = \frac{\sum_{i=1}^N \sum_{j=1}^{W_i} f(\mathbf{x}_{m,i})_j}{\sum_{i=1}^N W_i} \quad \text{and} \quad Y'_m = \frac{\sum_{i=1}^{N'} \sum_{j=1}^{W'_i} f(\mathbf{x}'_{m,i})_j}{\sum_{i=1}^{N'} W'_i}.$$

For the images in the style of writer m , the HWD is given by:

$$\text{HWD}_m = \|Y_m - Y'_m\|_2.$$

Note that when computed on robust image representations, *e.g.*, obtained from a backbone trained on a semantic prediction task, the Euclidean distance is highly predictive of the perceptual similarity between images [14]. Finally, the HWD on datasets containing images in the style of M different authors is obtained as

$$\text{HWD} = \frac{1}{M} \sum_{m=1}^M \text{HWD}_m.$$

The HWD score has the non-negativity, symmetry, triangular inequality properties, but it is not guaranteed that it exhibits the identity of indiscernible elements property since two non-identical images might have $\text{HWD} = 0$ due to their representation via the pre-trained backbone. This is a desirable characteristic for the HTG task, in which text images containing different texts and different backgrounds should be at low (or even zero) HWD if they are written in the same handwriting style.

4 Experimental Analysis

For our experimental analysis of the proposed HWD score, we consider images from a number of multi-author and single-author datasets. We compare the HWD score against the FID

score in the variant proposed in [23], which is the common approach adopted in Styled HTG. The comparison is performed along two dimensions. First, we assess its capability to recognize corresponding handwriting styles, quantifying the style verification capability with the Overlap coefficient and the Equal Error Rate (EER). Second, we compare the numerical stability of the proposed HWD score and the FID. Additionally, we evaluate current Styled HTG State-of-the-Art models with our score and other common metrics for image generation evaluation. Finally, we conduct extensive ablation analyses on the main components of our approach to investigate their individual contributions. Further results can be found in the Supplementary material.

4.1 Considered Datasets

The considered multi-author datasets are described below and in Table 1.

IAM. The IAM Database [30] is a collection of greyscale document scans by 657 writers. These are written in English with ink on white paper and cleaned digitally. Here, we use the line-level version of the dataset.

RIMES. The RIMES Database [0] consists of binary French documents by 1500 authors. For this dataset, we use the line-level version.

CVL. The CVL Database [26] features word images obtained from RGB scans of English and German manuscripts, written with ink on white paper by 310 writers.

KHATT. The KHATT Database [29] contains binarized images of handwritten Arabic words handwritten by 838 people.

BanglaWriting. The BanglaWriting Dataset [32] is composed of greyscale images of Bengalese words handwritten by 212 authors.

NorHand. The NorHand Dataset [28] features text lines extracted from greyscale scans of ancient documents written with ink on yellowed paper by 12 Norwegian authors. The considered single-author datasets are presented below.

Saint Gall. The Saint Gall Dataset [24] features binary images of 1410 lines from a medieval manuscript written in Latin with gothic calligraphy.

Washington. The George Washington Dataset [15] contains binary images of 656 lines from English letters written by American President George Washington.

Rodrigo. The Rodrigo Database [37] contains 20357 lines extracted from greyscale scans of an historical manuscript, written in Spanish with ink on ancient paper.

ICFHR14. The ICFHR 14 Dataset [36] is a collection of 11473 lines extracted from greyscale scans of ancient pages, written by the English philosopher Jeremy Bentham.

Leopardi. The Leopardi Dataset [8] is a collection of 2459 lines from RGB scans of letters by the Italian poet Giacomo Leopardi, written with ink on ancient paper.

LAM. The LAM Dataset [10] includes 25823 text lines images obtained from RGB scans of ancient letters in Italian, written by Lodovico Antonio Muratori.

4.2 Sensitivity to the Handwriting

We evaluate the sensitivity of HWD to the handwriting style by splitting the multi-author datasets. We consider half of the samples for each featured writer as references and the other half as if they were the output of a perfect Styled HTG model. Then, we compare the

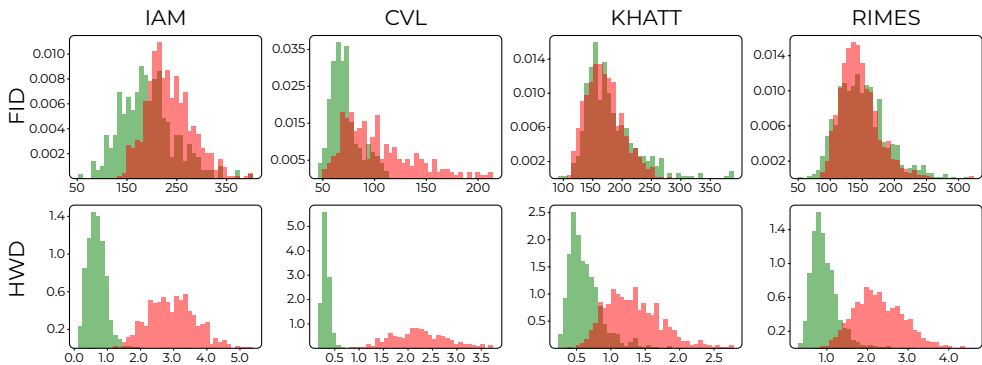


Figure 2: Distributions of different scores used to evaluate HTG models when applied on same-author (green) or different-author (red) subsets. The overlap area is in dark red.

distributions of the HWD and the FID computed on text images of multiple matching and non-matching authors pairs. Note that, in such an ideal case for Styled HTG, both the HWD and the FID should be as close as possible to their best value. Therefore, the more these two distributions are separated, the better the corresponding score captures the handwriting similarity between the considered images.

The obtained distributions are reported in Figure 2. We observe that the histograms for the FID show significant overlap and that there is no clear separation between the distributions of matching and non-matching authors pairs. Moreover, except for CVL, which has a higher number of samples per author compared to the other three datasets, the FID values of the corresponding authors’ distribution are roughly above 100. Such high FID values highlight the bias of the score when computed on a few samples. On the other hand, the histograms for the HWD are more separated.

In addition, we quantitatively evaluate the style recognition capability by comparing the distributions in Figure 2 in terms of the Overlap coefficient, a statistical index that quantifies the overlap between distributions, and the EER, which represents the point where the False Acceptance Rate equals the False Rejection Rate. The results are reported in Table 1. As we can see, for both the FID and HWD, the Overlap and ERR are lower on the IAM and CVL datasets than in KHATT and RIMES. We argue the causes to be, respectively, the languages of the datasets and the similar overall appearance of the samples from the same author. In fact, the FID mainly focuses on the latter, while HWD performs feature extraction with a VGG16 pretrained on Font², which contains words in the same language as IAM and CVL. Nonetheless, the HWD score achieves lower Overlap and EER values than the FID in all datasets, including those in languages different from English.

4.3 Sensitivity to the Number of Samples

Numerical stability is an important factor to consider when assessing a score. As argued by [2, 10], the FID exhibits a strong bias towards the number of samples. To assess the stability of HWD, we consider the large single-author LAM dataset and compute the values of the HWD and FID on images from LAM against variably sized subsets of images from ICFHR14, Saint Gall, Leopardi, Rodrigo, Washington, and LAM itself. We determine the mean of the scores over multiple runs and also consider the range between the 25th and 75th

| | Language | Samples | Authors | Avg. Samples per author | FID | | HWD | |
|----------------------|----------------|---------|---------|----------------------------|---------|------|---------|-----|
| | | | | | Overlap | EER | Overlap | EER |
| Norhand | Norwegian | 21939 | 12 | 1828.25 | 4.2 | 4.2 | 0.0 | 0.0 |
| BanglaWriting | Bengali | 17265 | 212 | 81.44 | 11.6 | 5.6 | 6.1 | 2.9 |
| CVL | English/German | 13473 | 310 | 43.46 | 24.7 | 12.5 | 0.0 | 0.0 |
| IAM | English | 13353 | 657 | 20.32 | 27.1 | 13.6 | 0.7 | 0.3 |
| KHATT | Arabic | 11427 | 838 | 13.64 | 40.3 | 21.6 | 12.0 | 5.9 |
| RIMES | French | 12111 | 1500 | 8.07 | 39.1 | 20.8 | 7.0 | 3.3 |

Table 1: Overlap and EER of the FID and HWD values calculated on the images in the considered multi-author datasets.

percentiles of values. The results are reported in Figure 3.

It can be observed that the FID (left-most plot) shows a significant bias towards the number of images: all the curves start with very high values and decrease slowly until at least around 2000 samples are used to obtain the score. On the other hand, the HWD (right-most plot) is more stable with respect to the subset sizes, reaching a plateau even when computed on around 100 samples. To further investigate the cause of this behavior, we compute the Euclidean distance on the Inception-v3 features (FID w/ Euclidean) and the Fréchet distance on the VGG16 features (HWD w/ Fréchet). In both cases (center-left plot and center-right plot of Figure 3, respectively), we can observe a slight bias with respect to the number of samples used for the computation. Moreover, the computation of these scores necessitates roughly one more order of magnitude of images to reach a plateau, compared to the case of the FID and the HWD, respectively. For the FID w/ Euclidean, we argue that this is because its Inception-v3 backbone is applied to the beginning part of the text images. As a result, more samples are needed to effectively represent the author’s handwriting. In the case of the HWD w/ Fréchet, we argue that this score suffers from the numerical instability of the Fréchet Distance, which fits the image representations to multivariate Gaussians.

Further, by comparing the values in the plots in the sight of the exemplar images reported in Figure 3, we can make some considerations on the convolutional backbones used as feature extractors. Inception-v3-based scores assign dataset distances according to their overall appearance. For instance, ICFHR14 and Rodrigo (both containing greyscale images) are close to each other, and the same happens for Saint Gall and Washington (both containing binarized images). On the other hand, the VGG16-based scores focus more on the handwriting. Thus, Saint Gall is isolated from the others because of its peculiar gothic calligraphy, while ICFHR14 is closer to Washington and Leopardi than to Rodrigo, reflecting the similarity in the cursive calligraphies. This characteristic is better suited for the Styled HTG task, as we are interested in evaluating the ability of a model to mimic the handwriting style.

4.4 Styled HTG Models Evaluation

For reference, we report the performance of State-of-the-Art Styled HTG models trained on the IAM dataset. In particular, we consider the conditional-GAN-based HIGAN+ approach [18], which uses a single image to condition the generation and disentangles text and style, and the few-shot Convolutional-Transformer-based models HWT [9] and VATr [62], which capture both global and character-level style-text dependencies by exploiting cross-attention. The latter approach represents the textual input with visual archetypes (*i.e.*, Unifont-rendered characters) instead of one-hot vectors, as done in the other two. In Table 2, we evaluate the aforementioned models in terms of the proposed HWD and other scores for

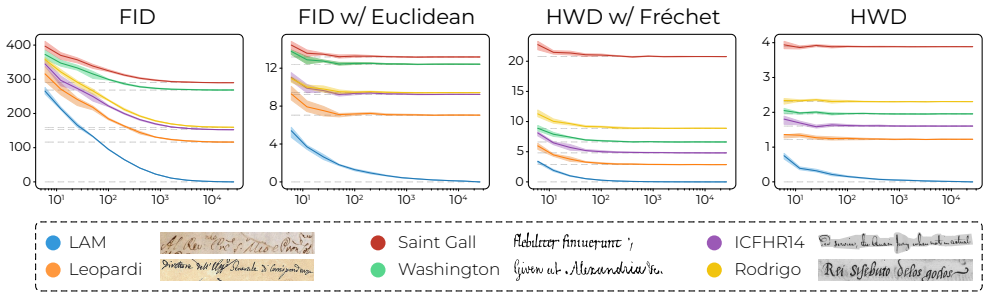


Figure 3: Left to right: comparison between FID, a version of the FID exploiting the Euclidean distance (FID w/ Euclidean), a version of HWD exploiting the Fréchet Distance (HWD w/ Fréchet), and HWD with varying number of samples. The lines denote the mean, and the transparent bands represent the range between the 25th and 75th percentiles, obtained with 10 calculation runs.

| | FID | KID | GS | HWD |
|--------|--------------|-----------------------|-----------------------|--------------|
| HWT | 23.36 | 1.37×10^{-2} | 1.05×10^{-2} | 1.928 |
| HIGAN+ | 18.21 | 9.38×10^{-3} | 2.15×10^{-2} | 1.237 |
| VATr | 18.80 | 7.06×10^{-3} | 2.19×10^{-2} | 0.828 |

Table 2: Scores of the considered HTG models when generating the same or different words as those in the style reference images of the IAM test set. Best performance in bold.

image generation evaluation. In particular, we consider the FID score; the Geometric Score (GS) [25], which compares the data and model manifold estimates; and the Kernel Inception Distance (KID) [7], which relaxes the Gaussian assumption of the FID and uses the Maximum Mean Discrepancy to compare the distributions of the features.

4.5 Ablation Analysis

Finally, we analyze the effects of the four main components of our proposed score: the backbone, the pretraining dataset, the input image portion, and the distance measure. We perform the ablation analysis on the IAM dataset and report the results in Table 3. Looking at the results, we notice that the backbone used to extract the features plays a crucial role in the separability scores, as the use of VGG16 leads to very good results also when pretrained on ImageNet. A second important aspect is the feature distance metric used. The Fréchet distance on the VGG16-extracted features achieves good results on all settings. Nevertheless, it is not influenced by the backbone pretraining and the image portion used to extract the features. On the other hand, the Euclidean distance fully exploits the input information (*i.e.*, the quantity and the type of feature vectors) and thus is the best in the HWD setting.

4.6 Sensitivity to the Visual Appearance

To assess the sensitivity to handwriting-related visual aspects, we compare the FID and HWD between reference images and increasingly altered ones, taken from the LAM dataset. In particular, the considered alterations entail shear, erosion, and dilation to simulate handwriting slant and strokes thickness. The results are reported in Figure 4 and show that HWD is more

| Backbone | Pretraining Dataset | Image Portion | Distance | IAM | |
|--------------|---------------------|---------------|-----------|------------|------------|
| | | | | Overlap | EER |
| Inception-v3 | ImageNet | Beginning | Fréchet | 27.1 | 13.6 |
| Inception-v3 | ImageNet | Beginning | Euclidean | 29.6 | 14.5 |
| Inception-v3 | ImageNet | Whole | Fréchet | 24.0 | 11.6 |
| Inception-v3 | ImageNet | Whole | Euclidean | 8.5 | 3.9 |
| Inception-v3 | Font ² | Beginning | Fréchet | 18.8 | 9.3 |
| Inception-v3 | Font ² | Beginning | Euclidean | 11.3 | 4.8 |
| Inception-v3 | Font ² | Whole | Fréchet | 19.0 | 9.1 |
| Inception-v3 | Font ² | Whole | Euclidean | 7.2 | 3.3 |
| VGG16 | ImageNet | Beginning | Fréchet | 3.2 | 1.6 |
| VGG16 | ImageNet | Beginning | Euclidean | 26.2 | 13.0 |
| VGG16 | ImageNet | Whole | Fréchet | 2.8 | 1.2 |
| VGG16 | ImageNet | Whole | Euclidean | 6.2 | 2.9 |
| VGG16 | Font ² | Beginning | Fréchet | 3.4 | 1.7 |
| VGG16 | Font ² | Beginning | Euclidean | 16.5 | 8.2 |
| VGG16 | Font ² | Whole | Fréchet | 3.5 | 1.6 |
| VGG16 | Font ² | Whole | Euclidean | 0.7 | 0.3 |

Table 3: Ablation analysis of the main components of the HWD score. Note that the first row is the FID while the last is the complete HWD. Best performance in bold.

sensitive than the FID to such visual aspects and that it increases linearly with the alteration intensity, thus enforcing its suitability for evaluating HTG.

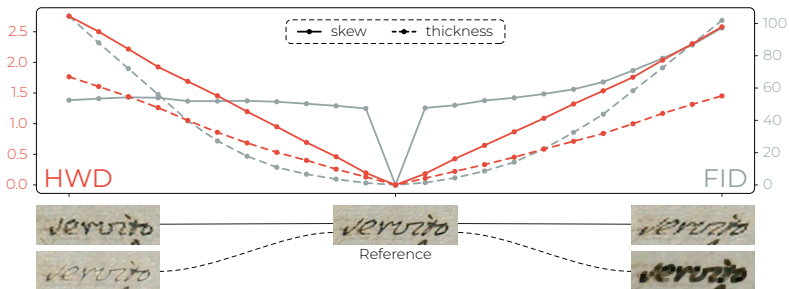


Figure 4: FID and HWD by varying handwriting thickness and slant.

5 Conclusion

In this work, we have proposed HWD, a score specifically designed to evaluate Styled HTG. HWD exploits the features extracted by a convolutional backbone trained on a large synthetic dataset of handwritten text to compare the perceptual differences between handwritings. Moreover, it is designed to work with images of variable lengths, such as those containing text. The results obtained from extensive experimental analysis demonstrate its suitability for evaluating text image generation approaches, its sensitivity to different styles, and its numerical stability. Hopefully, the use of the proposed score, whose implementation and backbone model weights will be made publicly available, will contribute to pushing forward the research on the Styled HTG task.

References

- [1] Eloi Alonso, Bastien Moysset, and Ronaldo Messina. Adversarial Generation of Handwritten Text Images Conditioned on Sequences. In *ICDAR*. IEEE Computer Society, 2019.
- [2] Emmanuel Augustin, Matthieu Carré, Emmanuèle Grosicki, J-M Brodin, Edouard Geoffrois, and Françoise Prêteux. RIMES evaluation campaign for handwritten mail processing. In *IWFHR*, 2006.
- [3] Ankan Kumar Bhunia, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Fahad Shahbaz Khan, and Mubarak Shah. Handwriting Transformers. In *ICCV*, 2021.
- [4] Ayan Kumar Bhunia, Abhirup Das, Ankan Kumar Bhunia, Perla Sai Raj Kishore, and Partha Pratim Roy. Handwriting Recognition in Low-Resource Scripts Using Adversarial Learning. In *CVPR*. IEEE, 2019.
- [5] Ayan Kumar Bhunia, Shuvojit Ghose, Amandeep Kumar, Pinaki Nath Chowdhury, Aneeshan Sain, and Yi-Zhe Song. MetaHTR: Towards Writer-Adaptive Handwritten Text Recognition. In *CVPR*, 2021.
- [6] Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, and Yi-Zhe Song. Text is Text, No Matter What: Unifying Text Recognition using Knowledge Distillation. In *ICCV*, 2021.
- [7] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- [8] Silvia Cascianelli, Marcella Cornia, Lorenzo Baraldi, Maria Ludovica Piazzini, Rosiana Schiuma, and Rita Cucchiara. Learning to Read L'Infinito: Handwritten Text Recognition with Synthetic Training Data. In *Proceedings of the International Conference on Computer Analysis of Images and Patterns*, 2021.
- [9] Silvia Cascianelli, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Boosting Modern and Historical Handwritten Text Recognition with Deformable Convolutions. *IJDAR*, 25(3):207–217, 2022.
- [10] Silvia Cascianelli, Vittorio Pippi, Maarand Martin, Marcella Cornia, Lorenzo Baraldi, Kermorvant Christopher, and Rita Cucchiara. The LAM Dataset: A Novel Benchmark for Line-Level Handwritten Text Recognition. In *ICPR*, 2022.
- [11] Min Jin Chong and David Forsyth. Effectively unbiased fid and inception score and where to find them. In *CVPR*, 2020.
- [12] Iulian Cojocaru, Silvia Cascianelli, Lorenzo Baraldi, Massimiliano Corsini, and Rita Cucchiara. Watch Your Strokes: Improving Handwritten Text Recognition with Deformable Convolutions. In *ICPR*. IEEE Computer Society, 2021.
- [13] Brian Davis, Chris Tensmeyer, Brian Price, Curtis Wigington, Bryan Morse, and Rajiv Jain. Text and Style Conditioned GAN for Generation of Offline Handwriting Lines. In *BMVC*, 2020.

- [14] Andreas Fischer, Volkmar Frinken, Alicia Fornés, and Horst Bunke. Transcription alignment of Latin manuscripts using hidden Markov models. In *HIP*, 2011.
- [15] Andreas Fischer, Andreas Keller, Volkmar Frinken, and Horst Bunke. Lexicon-free handwritten word spotting using character HMMs. *Pattern Recognition Letters*, 33(7): 934–942, 2012.
- [16] Sharon Fogel, Hadar Averbuch-Elor, Sarel Cohen, Shai Mazor, and Roei Litman. ScrabbleGAN: Semi-Supervised Varying Length Handwritten Text Generation. In *CVPR*, 2020.
- [17] Ji Gan and Weiqiang Wang. HiGAN: Handwriting Imitation Conditioned on Arbitrary-Length Texts and Disentangled Styles. In *AAAI*, 2021.
- [18] Ji Gan, Weiqiang Wang, Jiayu Leng, and Xinbo Gao. HiGAN+: Handwriting Imitation GAN with Disentangled Representations. *ACM Trans. Graphics*, 42(1):1–17, 2022.
- [19] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. Generative Adversarial Nets. In *NeurIPS*, 2014.
- [20] TSF Haines, O Mac Aodha, and GJ Brostow. My Text in Your Handwriting. *ACM Trans. Graphics*, 35(3), 2016.
- [21] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Advances in neural information processing systems*, 2017.
- [22] Lei Kang, Pau Riba, Marçal Rusinol, Alicia Fornes, and Mauricio Villegas. Distilling content from style for handwritten word recognition. In *ICFHR*, 2020.
- [23] Lei Kang, Pau Riba, Yaxing Wang, Marçal Rusiñol, Alicia Fornés, and Mauricio Villegas. GANwriting: Content-Conditioned Generation of Styled Handwritten Word Images. In *ECCV*, 2020.
- [24] Lei Kang, Pau Riba, Marçal Rusinol, Alicia Fornes, and Mauricio Villegas. Content and style aware generation of text-line images for handwriting recognition. *IEEE Trans. PAMI*, pages 1–1, 2021.
- [25] Valentin Khruikov and Ivan Oseledets. Geometry Score: A Method For Comparing Generative Adversarial Networks. In *ICML*. PMLR, 2018.
- [26] Florian Kleber, Stefan Fiel, Markus Diem, and Robert Sablatnig. CVL-DataBase: An Off-Line Database for Writer Retrieval, Writer Identification and Word Spotting. In *ICDAR*, 2013.
- [27] Canjie Luo, Yuanzhi Zhu, Lianwen Jin, Zhe Li, and Dezhi Peng. SLOGAN: Handwriting Style Synthesis for Arbitrary-Length and Out-of-Vocabulary Text. *IEEE Trans. Neural Netw. Learn. Syst.*, 2022.
- [28] Martin Maarand, Yngvil Beyer, Andre Kåsen, Knut T Fosseide, and Christopher Kermorvant. A comprehensive comparison of open-source libraries for handwritten text recognition in norwegian. In *DAS*, 2022.

- [29] Sabri A Mahmoud, Irfan Ahmad, Wasfi G Al-Khatib, Mohammad Alshayeb, Mohammad Tanvir Parvez, Volker Märgner, and Gernot A Fink. KHATT: An open Arabic offline handwritten text database. *Pattern Recognition*, 47:1096–1112, 2014.
- [30] U-V Marti and Horst Bunke. The IAM-database: an English sentence database for offline handwriting recognition. *IJDAR*, 5(1):39–46, 2002.
- [31] Alexander Mattick, Martin Mayr, Mathias Seuret, Andreas Maier, and Vincent Christlein. SmartPatch: Improving Handwritten Word Imitation with Patch Discriminators. In *ICDAR*, 2021.
- [32] Muhammad F Mridha, Abu Quwsar Ohi, M Ameer Ali, Mazedul Islam Emon, and Muhammad Mohsin Kabir. BanglaWriting: A multi-purpose offline Bangla handwriting dataset. *Data in Brief*, 34:106633, 2021.
- [33] Vittorio Pippi, Silvia Cascianelli, Lorenzo Baraldi, and Rita Cucchiara. Evaluating synthetic pre-training for handwriting processing tasks. *Pattern Recognition Letters*, 2023.
- [34] Vittorio Pippi, Silvia Cascianelli, and Rita Cucchiara. Handwritten Text Generation from Visual Archetypes. In *CVPR*, 2023.
- [35] Vittorio Pippi, Silvia Cascianelli, Christofer Kermorvant, and Rita Cucchiara. How to Choose Pretrained Handwriting Recognition Models for Single Writer Fine-Tuning. In *ICDAR*, 2023.
- [36] Joan Andreu Sánchez, Verónica Romero, Alejandro H Toselli, and Enrique Vidal. ICFHR2014 competition on handwritten text recognition on transcriptorium datasets (HTRtS). In *ICFHR*, 2014.
- [37] Nicolás Serrano, Francisco Castro, and Alfons Juan. The RODRIGO Database. In *LREC*, 2010.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [39] Jue Wang, Chenyu Wu, Ying-Qing Xu, and Heung-Yeung Shum. Combining Shape and Physical Models for On-line Cursive Handwriting Synthesis. *IJDAR*, 7(4):219–227, 2005.
- [40] Pei Wang, Yijun Li, and Nuno Vasconcelos. Rethinking and Improving the Robustness of Image Style Transfer. In *CVPR*, 2021.
- [41] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [42] Yaping Zhang, Shuai Nie, Wenju Liu, Xing Xu, Dongxiang Zhang, and Heng Tao Shen. Sequence-to-sequence domain adaptation network for robust text image recognition. In *CVPR*, 2019.