HR-FONT: CROSS-SCRIPT FONT GENERATION VIA REALIZATION PRIORS AND APPEARANCE COMPLETION

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce HR-Font, a framework for extending a typeface to new scripts from a few Chinese reference glyphs. The central challenge is to translate a family's visual language into characters whose structures differ from those observed. HR-Font addresses this challenge through two complementary forms of conditioning. The Character Realization Prior (CRP) represents same-character examples from a font bank as variations relative to neutral content and uses them to condition spatial deformation in the denoiser. Target-character appearance completion reads multi-scale local reference features to predict a spatial appearance representation of the requested glyph, supervised by training-target features. The resulting conditions guide a diffusion generator toward recognizable characters with a coherent family identity. We evaluate paired reconstruction and cross-script reference-style consistency on 47 fonts across Latin, kana, and Bopomofo. In a blinded Latin-character study with 28 designers and domain experts, HR-Font is the highest-ranked generated method, with a 77.5% Top-2 inclusion rate.

1 INTRODUCTION

Cross-script font generation synthesizes glyphs in one writing system from style references in another. We study the setting in which a few Chinese glyphs serve as style references for generating Latin, kana, and Bopomofo characters. The generated glyph must preserve character identity and reflect the references' stroke treatment and local appearance. Chinese references reveal properties such as stroke weight, contrast, contour treatment, and decoration, but expressing them in a target script requires character-specific choices of geometry and appearance.

Few-shot font generators commonly combine a content glyph with reference-style features. Prior work develops localized and fine-grained style representations (Park et al., 2021a;b; Tang et al., 2022; Pan et al., 2023), and enriches content representations with examples from basis fonts (Wang et al., 2023). Recent work further explores diffusion-based font synthesis, font-transfer processes, autoregressive generation, and explicit glyph spatial modeling (Yang et al., 2024; Fu et al., 2024; 2025; Cai et al., 2026; Su & Yang, 2026). Cross-script methods further study reference attention and multilingual transfer (Li et al., 2021; Zhang et al., 2022).

However, in the cross-script setting, the source references leave two aspects underdetermined, as illustrated in Figure 1: how the requested target character should be geometrically realized, and how source-script appearance should be expressed on target-character structures.

First, character identity permits substantial geometric variation. Typefaces render the same character with different proportions, stroke placements, curves, and counter shapes. A neutral content glyph gives one such realization. At inference, the Chinese references provide no example of the requested target character in the desired style. Realizations from other fonts define a character-specific prior over plausible structural variations around the neutral glyph, while the observed references guide which variations are more compatible with the target style.

Local appearance must also be adapted to the target character. The terminals, junctions, and curves of a target character may have no direct counterparts in the Chinese references. The model must determine how local appearance cues observed in the references should be expressed at target-character locations with different structural roles. This motivates a target-aligned representation that

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

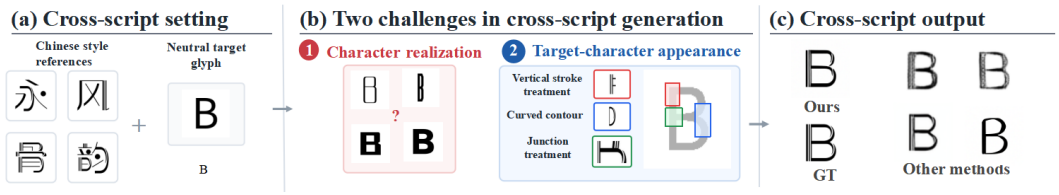


Figure 1: Overview of the cross-script font generation setting and its two key challenges. Given a neutral target glyph and a few Chinese reference glyphs, the model must determine a plausible target-character realization and adapt source-script appearance cues to the target-character structure.

aggregates local observations from the references and is trained to predict spatial features of the target glyph.

We propose HR-Font, a diffusion-based model with two target-character conditioning modules. The *Character Realization Prior* (CRP) first matches the Chinese references to renderings of the same characters in candidate training fonts, then obtains renderings of the requested target character from the selected fonts. It represents these candidate realizations as feature differences from the neutral target through parallel content- and style-encoder branches. Each branch dynamically routes the differences and predicts spatial offsets. The predicted offsets are fused and applied through a deformable residual update of the denoiser's skip features. The selected candidates remain fixed during sampling; routing adapts to each location, structural module, and denoising step.

Target-Character Appearance Completion (TC) independently encodes each Chinese reference at two scales. Queries derived from the neutral target attend to the reference tokens to form a target-aligned representation. Two projections of this shared representation produce local appearance tokens for the denoiser and a spatial prediction supervised by features of the paired training target.

CRP supplies structural offsets from candidate realizations, while TC predicts appearance features on a target-content grid. The two conditions guide a FontDiffuser-based generator alongside its content and global-style features. We evaluate paired reconstruction and cross-script reference-style consistency using matched reference sets and controlled reference budgets.

A paired target provides one concrete realization, while source-script references can support other plausible designs. Reconstruction metrics are therefore complemented by an independently trained cross-script style evaluator that compares each generated target glyph with its Chinese references.

Our contributions are:

- We develop HR-Font for cross-script font generation, conditioning a diffusion model on target-character structural variations and a target-aligned appearance representation. In human ranking study, HR-Font receives the highest overall ranking among generated methods.
- We introduce CRP, which represents reference-selected realizations of the same target character relative to neutral content in two encoder spaces and uses them to condition skip-feature deformation.
- We introduce TC, which predicts target-character spatial features from Chinese references and uses the shared representation to provide local appearance conditions for generation.

2 Related Work

2.1 Few-Shot Font Generation

Few-shot font generation (FFG) synthesizes a typeface from a few reference glyphs, typically by disentangling character content from typeface style. LF-Font models localized component-wise style representations, while MX-Font uses multiple localized experts to capture reusable local patterns (Park et al., 2021a;b). FSFont and VQ-Font further model fine-grained local style transfer through spatial correspondence and learned component-level representations (Tang et al., 2022; Pan et al., 2023). These methods establish the importance of localized representations for preserving detailed font appearance.

Another line of work focuses on content representation and generative modeling. CF-Font enriches content features using corresponding glyphs from multiple basis fonts (Wang et al., 2023). FontDiffuser and subsequent diffusion-based methods model font synthesis through conditional denoising and font-transfer processes (Yang et al., 2024; Fu et al., 2024; 2025). More recent work explores global-aware autoregressive generation and explicit modeling of glyph spatial information (Cai et al., 2026; Su & Yang, 2026). These methods improve content fidelity, structural quality, and local detail in few-shot generation. HR-Font considers the related but different setting in which the structural regularities of the reference script do not directly determine the realization of a target-script character.

2.2 Cross-Script Font Generation

Cross-script font generation transfers typeface characteristics between writing systems with different character structures. FTransGAN explicitly studies few-shot transfer between languages and uses multi-level attention to capture global and local reference information (Li et al., 2021). MF-Net extends multilingual font generation to previously unseen styles and languages using separate content and style encoders together with language-aware structural connections (Zhang et al., 2022). FCA-GAN improves cross-lingual style extraction through full-domain convolutional attention and multi-level feature modeling (Zhao et al., 2024), while DAA-Font combines dual-branch encoding with multi-level axial attention to model style and structural information across scripts (Yan et al., 2026). RPFont further improves cross-lingual transfer through patch-level style contrastive learning and relative-position-aware modeling (Zeng et al., 2026). These methods improve structural fidelity and style preservation when the target character is specified by a content glyph. In cross-script generation, however, the source references do not reveal which geometric realization of that character is preferred by the target style. Our Character Realization Prior (CRP) represents these font-dependent geometric variations relative to a neutral realization and provides them as a target-specific structural condition. This separates the variation of target geometry from the appearance information inferred from the source-script references.

2.3 Font Evaluation and Style Representation

Font generation is commonly evaluated with reconstruction, perceptual, distributional, and recognition-based measures, including SSIM (Wang et al., 2004), LPIPS (Zhang et al., 2018), FID (Heusel et al., 2017), and content/style classifiers (Park et al., 2021a). These measures characterize different aspects of generation quality, but most compare images within a common script or against a particular target realization.

Learned font-style representations provide another way to measure typeface similarity. Hassan et al. learn a font-style embedding space with metric learning, placing glyphs with similar font styles closer in the representation space (Hassan et al., 2024). Relatedly, RPFont uses patch-level contrastive learning to learn stylistic representations that transfer across languages (Zeng et al., 2026). These works show that font appearance can be compared in a learned representation space rather than only through image-level similarity. For cross-script evaluation, we train an independent evaluator that learns a shared typeface representation across scripts and measures the consistency between source-script references and generated target glyphs.

3 Method

3.1 Overview

cross-script font generation requires a target-character realization whose structure and local details are consistent with the reference family. Chinese references reveal stroke treatment and decoration, while the neutral target provides one possible geometry. Let $x_{f,a}$ denote the rendering of character a in typeface f. Given Chinese references $\mathcal{R}_f = \{r_i = x_{f,a_i}\}_{i=1}^{n}$ and a neutral rendering x_c^0 of target character c, we generate the target glyph in family f. We construct two target-character conditions. The *Character Realization Prior* (CRP) uses reference similarity to select candidate training families, then routes feature differences between their target-character renderings and the neutral glyph. Two parallel branches convert these differences into structural offsets. *Target-Character*

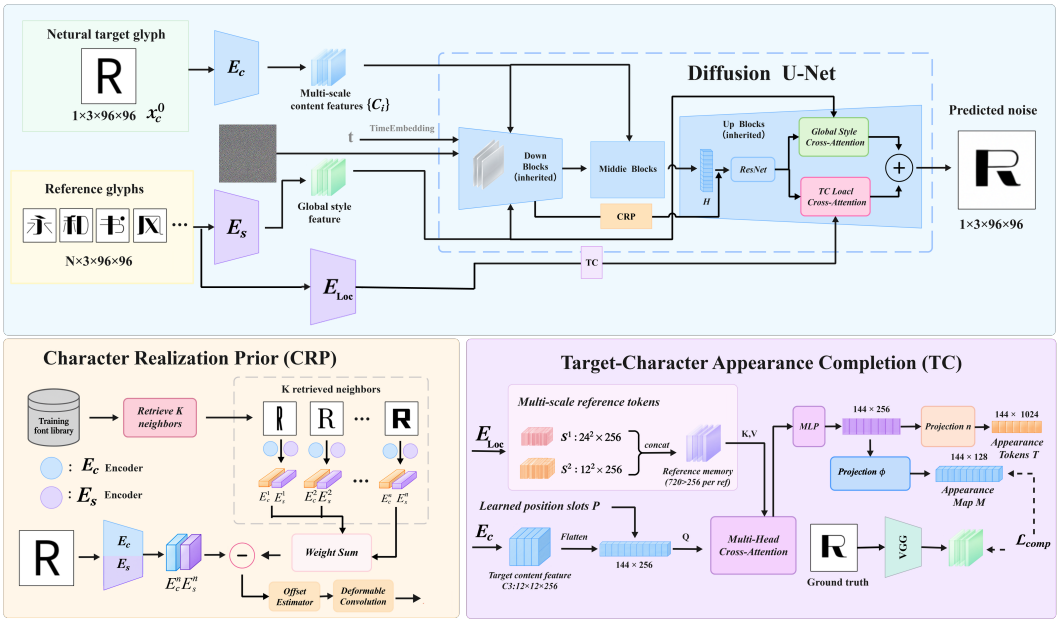


Figure 2: HR-Font overview. CRP supplies structural offsets to skip features, and TC supplies local appearance tokens to decoder cross-attention.

Appearance Completion (TC) predicts target-aligned local features from the references and supplies local appearance tokens to the denoiser.

We use a FontDiffuser-based denoiser (Yang et al., 2024), retaining its multi-scale content aggregation (MCA) and global-style conditioning. A frozen content encoder E_c provides target-content features, and a frozen style encoder E_s provides global reference tokens G and descriptors for candidate selection. Before HRFont training, we fine-tune the inherited E_s and E_c on the cleaned training split; we then freeze them and reuse matching cached features, computing missing features on demand. TC uses a separate trainable local encoder. CRP modifies skip features through spatial offsets, while TC conditions decoder cross-attention. Figure 2 summarizes the model.

3.2 Character Realization Prior

Reference-guided candidates. We compare the observed references with renderings of the same Chinese characters in a fixed pool of training families. An admissible family contains the requested target character and all observed reference characters. The episode font and its explicit weight variants are excluded before selection. Using the pooled descriptor g from E_s, we compute

$$s_j = \frac{1}{n}\sum_{i=1}^{n} \cos\big(g(r_i), g(x_{f_j,a_i})\big)\,,\ \alpha_j = \frac{\exp(s_j/\tau)}{\sum_{k\in\mathcal{J}}\exp(s_k/\tau)},\ j \in \mathcal{J}. \tag{1}$$

Here $\mathcal{J} = \mathcal{J}(c, \mathcal{R}_f)$ contains up to K highest-scoring admissible families, and $\tau > 0$ is the retrieval temperature. Each selected family provides a rendering $x_{f_j,c}$ of the requested target character. The candidate set and prior α remain fixed throughout denoising.

Parallel feature differences. The two branches represent each candidate relative to the neutral target using spatial features from the frozen encoders:

$$d^{\mathrm{ec},s}_{j,c} = E^s_c(x_{f_j,c}) - E^s_c(x^0_c),\ d^{\mathrm{es},s}_{j,c} = E^s_s(x_{f_j,c}) - E^s_s(x^0_c). \tag{2}$$

Here s indexes the feature level. At each structural module, the E_s difference passes through a bias-free 1×1 adapter. The branches share the candidates and prior weights, with separate routers and offset predictors.

Candidate aggregation and structural offsets. CRP maps candidate feature differences to spatial offsets for skip-feature deformation. In our implementation, candidate aggregation uses normalized, state-dependent weights within each branch.

Let $U_{\ell,t}$ be the skip feature at structural module ℓ and denoising step t. Write $\bar{d}^b_{j,\ell}$ for the candidate difference at the matching feature level, including the adapter in the E_s branch. Each branch forms a query $q^b_{\ell,t,p}$ at position p from the skip feature, the neutral-content feature, the global reference summary, and the timestep. For $b \in \{\mathrm{ec}, \mathrm{es}\}$, we compute

$$w^b_{j,\ell,t,p} = \mathrm{softmax}_{j\in\mathcal{J}}\left[\log \alpha_j + \gamma^b_\ell \cos\big(q^b_{\ell,t,p}, K^b_\ell(\bar{d}^b_{j,\ell})_p\big)\right], \quad D^b_{\ell,t,p} = \sum_{j\in\mathcal{J}} w^b_{j,\ell,t,p}\bar{d}^b_{j,\ell,p},$$
$$\Delta^b_{\ell,t} = F^b_\ell(U_{\ell,t}, D^b_{\ell,t}), \quad \Delta_{\ell,t} = \Delta^{\mathrm{ec}}_{\ell,t} + \tanh(\zeta_\ell)\Delta^{\mathrm{es}}_{\ell,t}. \tag{3}$$

K^b_ℓ is a key projection, F^b_ℓ is an offset predictor, and γ^b_ℓ and ζ_ℓ are learned scalars.

The fused offset controls the residual skip update:

$$\widetilde{U}_{\ell,t} = U_{\ell,t} + P_\ell[\mathrm{DCN}_\ell(U_{\ell,t}, \Delta_{\ell,t}) - U_{\ell,t}]. \tag{4}$$

P_ℓ is a zero-initialized 1×1 projection. The updated skip is concatenated with the decoder feature before the corresponding residual block.

3.3 Target-Character Appearance Completion

TC predicts local appearance features at target-character positions from the neutral content and Chinese references. A shared representation produces a spatial feature prediction supervised by the paired target and a separate token representation for the denoiser.

A shared trainable encoder E_{loc} extracts a shallow and a deeper feature map from each reference independently. The shallow map is projected to the deeper map's channel width before both are flattened and concatenated into a token sequence S_i. Target-content queries jointly read the reference sequences:

$$\mathcal{S} = [S_1; \ldots; S_n], \quad Q_c = \mathcal{Q}(C_c) + \Pi_{\mathrm{slot}}, \quad Z_0 = \mathrm{Attn}(Q_c, \mathcal{K}(\mathcal{S}), \mathcal{V}(\mathcal{S})),$$
$$Z = Z_0 + \mathrm{FFN}(\mathrm{LN}(Z_0)), \quad M = \phi(Z), \quad T = \eta(Z). \tag{5}$$

C_c contains flattened target-content features from $E_c(x^0_c)$, Π_{slot} contains learned target-position vectors, and $\mathcal{Q}$, $\mathcal{K}$, and $\mathcal{V}$ are affine projections. Target alignment refers to the spatial grid indexed by the neutral target-content queries. Attention uses head-wise query/key normalization. For each query and head, its softmax spans all valid tokens across references, scales, and spatial positions. Independent affine projections ϕ and η produce the spatial prediction M and local appearance tokens T.

We supervise M with spatial VGG features (Simonyan & Zisserman, 2015) of the paired target $y = x_{f,c}$:

$$M^* = \mathrm{sg}[\mathcal{N}(\mathrm{Pool}(V_2(y)))], \quad \mathcal{L}_{\mathrm{comp}} = \mathrm{SmoothL1}(M, M^*). \tag{6}$$

V_2 denotes the frozen VGG feature stage used for local supervision. Pooling matches the target-query grid and flattens it into rows, $\mathcal{N}$ applies fixed channel-wise normalization, and sg stops gradients through the target features.

Decoder conditioning. At each decoder cross-attention module, global reference tokens G and local appearance tokens T are read with shared attention parameters and separate softmax normalizations. For an active reference condition,

$$A^G_{\ell,t} = \mathrm{CA}_\ell(H_{\ell,t}, G), \; A^T_{\ell,t} = \mathrm{CA}_\ell(H_{\ell,t}, T), \; \widetilde{A}_{\ell,t} = A^G_{\ell,t} + \beta_0 r(u)\kappa_{\ell,t}A^T_{\ell,t}. \tag{7}$$

where $H_{\ell,t}$ is the current attention input. The factor $\kappa_{\ell,t}$ is the detached, clipped global-to-local ratio of the outputs' root-mean-square (RMS) magnitudes. β_0 is a fixed coefficient, and $r(u)$ ramps the local contribution over training updates u.

3.4 Training and Inference

Optimization. Starting from a pretrained denoiser, we jointly train it with both structural branches and TC. The encoders E_c, E_s, and VGG remain frozen. Each episode pairs one to eight Chinese references with a target glyph from the same family. The objective is

$$\mathcal{L} = \mathcal{L}_{\text{diff}} + 0.01\mathcal{L}_{\text{perc}} + r(u)(0.01\mathcal{L}_{\text{comp}} + 0.05\mathcal{L}_{\text{detail}})\,. \tag{8}$$

$\mathcal{L}_{\text{diff}}$ is the mean squared error of the predicted noise. The perceptual loss compares VGG features of the paired target and the clean-image estimate reconstructed from the noisy input and predicted noise. The detail loss combines errors in target-defined regions, regions where the design adds or removes ink relative to the neutral rendering, and local high-frequency structure. Appendices A and D summarize the implementation and evaluation protocols.

Inference. The candidate set, prior, and TC tokens are computed once per content-reference pair and reused throughout sampling. CRP routing, decoder attention, and RMS calibration are recomputed at each denoising step. The ramp is set to one at inference. The denoiser predicts noise, and diffusion sampling produces the final glyph.

4 Experiments

4.1 Experimental setup

Data. The training set contains 137,639 font–character pairs from 738 fonts after filtering. The validation set contains 7,331 pairs from 32 fonts. The test set assesses whether distinctive styles are preserved across scripts and contains 8,638 valid pairs from 47 fonts and 284 target characters: 3,542 Latin, 4,726 kana, and 370 Bopomofo examples. The evaluation fonts are held out by font identity, while all target characters occur in the generator training data. Chinese glyphs provide the style references. Images are rendered at 96×96 with a shared font size within each font, and neutral content uses Noto Sans CJK Regular.

Settings. We reproduce or fine-tune FTransGAN, FCA-GAN, GAR-Font, and FontDiffuser on the same training data, using their original generation reference budgets (Table 1). HR-Font is trained on eight NVIDIA V100 GPUs with a global batch size of 64, and evaluation uses EMA weights. Checkpoints are selected from loss and L1 diagnostics together with visual assessment of validation outputs. Model configurations and reference protocols are given in Appendix D.

Metrics. We report *Style Compatibility* (SC), the cosine similarity between a generated glyph and a Chinese-reference style prototype in an independently trained, frozen font encoder.The evaluator is trained separately from the generator, with evaluation fonts excluded from its training set. The evaluator is used after checkpoint selection and does not participate in generator training or model selection. SC does not require the target character's ground-truth rendering, and its evaluation references are fixed across methods and generation budgets. AlexNet-based LPIPS (Zhang et al., 2018), L1, and SSIM (Wang et al., 2004) measure agreement with the paired target rendering. Higher SC and SSIM and lower LPIPS and L1 indicate better scores. Automatic metrics are averaged over characters within each font, then over fonts with equal weight. Paired comparisons use font-level resampling to obtain 95% confidence intervals.

4.2 Comparison with existing methods

Table 1 reports each method's reference budget and includes HR-Font with one and eight references. At the matched one-reference budget, HR-Font has higher SC than FontDiffuser, with a paired difference of 0.049523 and a 95% interval of $[0.023096,\ 0.082772]$. This supports improved reference-style compatibility in the single-reference setting. With eight references, HR-Font achieves the highest SC and lowest LPIPS among the reported configurations, reducing LPIPS by 5.3% relative to single-reference FontDiffuser. At one reference, FontDiffuser has lower L1 and LPIPS and higher SSIM. FCA-GAN obtains the lowest L1 and highest SSIM. Figure 3 compares the generated glyphs.

Table 1: Cross-script generation on the constructed test set. Ref. is the number of style references supplied to the generator. Values are font-macro averages; bold marks the best value in each column.

Method	Ref.	SC ↑	LPIPS ↓	L1 ↓	SSIM ↑
FTransGAN	6	0.480683	0.229734	0.063586	0.782643
FCA-GAN	6	0.475116	0.228859	**0.063130**	**0.783203**
GAR-Font	8	0.584269	0.205034	0.071588	0.767646
FontDiffuser	1	0.613962	0.186309	0.082526	0.741735
HR-Font	1	0.663486	0.191209	0.087498	0.733956
HR-Font	8	**0.699777**	**0.176435**	0.084549	0.739638



Figure 3: Main comparison for cross-script font generation. Chinese references and ground-truth renderings provide the design context for comparing the generated target glyphs.

Expert ranking study. Twenty-eight designers and domain experts each ranked six candidates for 30 Latin-character items, yielding 840 complete rankings. A separate expert group curated the items from ground-truth glyphs without inspecting generated outputs, selecting distinctive typefaces and characters that express their stylistic features. The selectors and raters did not overlap, and neither group included authors or model developers. Method identities and the ground-truth label were concealed, and candidate positions were randomized. Each item presents Chinese references, outputs from the five methods, and the ground-truth glyph as a sixth candidate. The HR-Font candidate uses eight generation references; all other candidates use the same model versions and reference budgets as Table 1. The task asks which candidates best match the reference typeface in stroke weight, proportions, contour treatment, and decoration. The questionnaire includes practice items and two style-mismatch attention checks; only formal rankings enter the results. We report mean rank and Top-2 inclusion rate, with 95% intervals obtained by bootstrapping items. The questionnaire and quality-control rules are given in Appendix F.

Experts rank HR-Font highest among the generated candidates by both measures (Table 2). Its Top-2 inclusion rate is 77.5%, compared with 22.5% for the best-ranked baseline and 80.5% for GT. These rankings favor HR-Font for reference-style consistency on the expert-curated Latin panel. Mean within-item agreement is Kendall's $W = 0.636$.

Metric agreement. Within each questionnaire item, we compare the ordering of generated candidate pairs under each automatic metric with the ordering of their expert mean ranks, then average agreement across items. GT is excluded from this analysis. SC and LPIPS agree with expert orderings in 66.8% and 74.0% of comparisons, respectively, compared with 43.2% for L1 and 47.9% for SSIM (Table 3). The SC and LPIPS intervals remain above 50% under both item-level and font-clustered resampling. These results support SC as a reference-conditioned measure of expert style preferences on this panel.

Table 2: Expert rankings of six candidates. Top-2 is the percentage of complete rankings placing a candidate first or second; brackets give 95% item-bootstrap intervals. Bold identifies the best generated candidate.

Candidate	Mean rank ↓	Top-2 (%) ↑
FTransGAN	4.782	3.5 [1.8, 5.4]
FCA-GAN	4.757	2.4 [1.2, 3.8]
GAR-Font	3.927	13.7 [8.8, 19.5]
FontDiffuser	3.600	22.5 [14.3, 31.7]
HR-Font	**2.111**	**77.5** [70.0, 84.3]
Ground truth	1.823	80.5 [72.7, 87.7]

Table 3: Agreement with expert within-item rankings (%). Brackets denote 95% confidence intervals.

	SC	LPIPS	L1	SSIM
Item bootstrap	66.8 [60.4, 73.0]	74.0 [67.3, 80.3]	43.2 [36.1, 50.3]	47.9 [40.0, 55.9]
Font-clustered	66.8 [60.5, 73.2]	74.0 [66.3, 81.1]	43.2 [36.3, 49.2]	47.9 [39.6, 55.6]

4.3 Ablation studies

We compare separately trained module configurations and inference-time candidate interventions. All comparisons use matched reference episodes, the same sampler, and fixed initial diffusion noise.

Structural and appearance conditions. We train Baseline, Baseline + CRP, Baseline + TC, and Full as separate configurations. Baseline omits both CRP and TC while retaining ordinary skip connections and global reference conditioning; Full denotes HR-Font with both modules. Full improves all four reported metrics over the other trained configurations (Table 4). Its LPIPS is 8.0% lower than Baseline + CRP, 12.1% lower than Baseline + TC, and 15.8% lower than Baseline. All paired metric differences favor Full, with 95% font-level bootstrap intervals above zero (Table 6).

Table 4: Module ablations and candidate interventions with eight references. The upper block contains separately trained configurations; the lower block changes candidate use in the fixed Full model. On and Off indicate whether each module is enabled. Full is HR-Font.

	Modules		Metrics			
Configuration	CRP	TC	SC ↑	LPIPS ↓	L1 ↓	SSIM ↑
Separately trained configurations						
Baseline	Off	Off	0.657041	0.209599	0.092786	0.721680
Baseline + CRP	On	Off	0.679041	0.191694	0.087784	0.732390
Baseline + TC	Off	On	0.676501	0.200755	0.091563	0.724683
Full	On	On	0.699777	0.176435	0.084549	0.739638
Inference-time interventions on Full						
Uniform Top-K	On	On	0.699714	0.176462	0.084568	0.739591
Retrieval Top-1	On	On	0.698470	0.171600	0.080835	0.748479
Random donors	On	On	0.695751	0.180701	0.087100	0.733961

Candidate selection and combination. With the trained Full model fixed, we replace learned candidate weights with uniform weights, keep only the highest-ranked candidate, or swap the retrieved candidates for an equally sized random set from the admissible pool. Reference-guided selection outperforms random donors, reducing LPIPS from 0.180701 to 0.176435. Uniform weighting closely matches the learned combination. Top-1 lowers L1 and LPIPS and raises SSIM, with SC of 0.698470 versus 0.699777 for the learned combination. Donor selection changes the metrics more than switching between learned and uniform weights. Full retains its performance under these alternative combination rules (Figure 4).

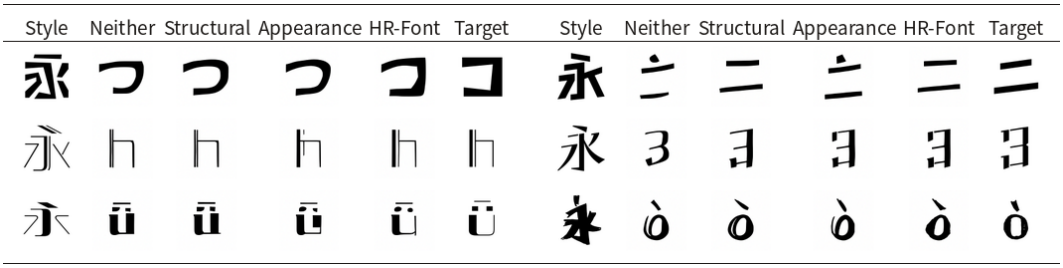


Figure 4: Visual comparison of separately trained module configurations and candidate interventions on the trained Full model. Target characters and reference glyphs are matched across conditions; Full denotes HR-Font.

4.4 Reference budgets

We evaluate one HR-Font checkpoint on strictly nested sets of one, two, four, and eight generation references, with targets and initial diffusion noise fixed. Each smaller set is a subset of the larger ones, and retrieval is recomputed from the references available at each budget. Additional references improve perceptual reconstruction and style compatibility (Table 5): LPIPS falls from 0.191209 to 0.176435 and SC rises from 0.663486 to 0.699777 between one and eight references. The largest LPIPS and SC changes occur between one and two references, with smaller subsequent gains. L1 and SSIM peak at four references, while LPIPS and SC continue to improve at eight. All settings use the same evaluation references.

Table 5: Reference-budget evaluation of HR-Font. The eight-reference row is the full-model configuration.

References	SC ↑	LPIPS ↓	L1 ↓	SSIM ↑
1	0.663486	0.191209	0.087498	0.733956
2	0.688275	0.180277	0.084951	0.739869
4	0.696130	0.179004	**0.084269**	**0.740546**
8	**0.699777**	**0.176435**	0.084549	0.739638

5 Discussion and Conclusion

We present HR-Font for cross-script font generation from Chinese reference glyphs. CRP represents candidate realizations of the target character relative to the neutral glyph, and TC predicts target-aligned appearance features from the references. Among the four separately trained module configurations, the full model achieves the highest SC and lowest LPIPS. With eight references, HR-Font also attains lower LPIPS than the compared baselines and ranks highest among generated methods in the Latin expert study.

SC and LPIPS agree more closely with expert within-item orderings than L1 and SSIM (Table 3). Uniform weighting closely matches learned weighting, while Top-1 improves paired reconstruction at similar SC (Table 4). Together with the module comparisons, these results suggest that the structural condition remains useful across the tested aggregation rules.

HR-Font currently generates isolated 96×96 raster glyphs. Quantitative evaluation covers Latin, kana, and Bopomofo; the expert study covers Latin. Future work includes higher-resolution or vector glyphs and inter-glyph spacing.

Reproducibility statement

Implementation, data preparation, and evaluation details are described in the paper and appendix. We will release the core training and inference code, model definitions, training configurations, evaluation data permitted for redistribution, and a versioned inventory of training fonts. Trained HR-Font checkpoints and the full internal codebase are not included in this release.

ETHICS STATEMENT

The expert evaluation involved 28 designers and domain experts who voluntarily ranked generated glyphs for consistency with Chinese reference typefaces. Any released evaluation records will exclude identifying information. Font assets and rendered glyphs will be redistributed only where permitted by their licenses or explicit authorization; otherwise, we will provide source and license information.

AI USE STATEMENT

AI tools assisted with related work discovery, manuscript editing and polishing. The authors take responsibility for the final text, claims, and results.

REFERENCES

Haonan Cai, Yuxuan Luo, and Zhouhui Lian. Beyond patches: Global-aware autoregressive model for multimodal few-shot font generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 517–528, 2026.

Bin Fu, Fanghua Yu, Anran Liu, Zixuan Wang, Jie Wen, Junjun He, and Yu Qiao. Generate like experts: Multi-stage font generation by incorporating font transfer process into diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6892–6901, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/html/Fu_Generate_Like_Experts_Multi-Stage_Font_Generation_by_Incorporating_Font_Transfer_CVPR_2024_paper.html.

Bin Fu, Zixuan Wang, Kainan Yan, Shitian Zhao, Qi Qin, Jie Wen, Junjun He, and Peng Gao. FontAnimate: High quality few-shot font generation via animating font transfer process. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16015–16025, 2025. URL https://openaccess.thecvf.com/content/ICCV2025/html/Fu_FontAnimate_High_Quality_Few-shot_Font_Generation_via_Animating_Font_Transfer_ICCV_2025_paper.html.

Ammar Ul Hassan, Irfanullah Memon, and Jaeyoung Choi. Learning font-style space using style-guided discriminator for few-shot font generation. *Expert Systems with Applications*, 242:122817, 2024. doi: 10.1016/j.eswa.2023.122817.

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*, volume 30, 2017.

Chenhao Li, Yuta Taniguchi, Min Lu, and Shin'ichi Konomi. Few-shot font style transfer between different languages. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 433–442, 2021.

Wei Pan, Anna Zhu, Xinyu Zhou, Brian Kenji Iwana, and Shilin Li. Few shot font generation via transferring similarity guided global style and quantization local style. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 19506–19516, 2023. URL https://openaccess.thecvf.com/content/ICCV2023/html/Pan_Few_Shot_Font_Generation_Via_Transferring_Similarity_Guided_Global_Style_ICCV_2023_paper.html.

Song Park, Sanghyuk Chun, Junbum Cha, Bado Lee, and Hyunjung Shim. Few-shot font generation with localized style representations and factorization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 2393–2402, 2021a. doi: 10.1609/aaai.v35i3.16340.

Song Park, Sanghyuk Chun, Junbum Cha, Bado Lee, and Hyunjung Shim. Multiple heads are better than one: Few-shot font generation with multiple localized experts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13900–13909, 2021b.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1409.1556.

Peng Su and Xi Yang. Rethinking glyph spatial information in font generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 29190–29199, 2026.

Licheng Tang, Yiyang Cai, Jiaming Liu, Zhibin Hong, Mingming Gong, Minhu Fan, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. Few-shot font generation by learning fine-grained local styles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7895–7904, 2022.

Chi Wang, Min Zhou, Tiezheng Ge, Yuning Jiang, Hujun Bao, and Weiwei Xu. CF-Font: Content fusion for few-shot font generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1858–1867, 2023.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. doi: 10.1109/TIP.2003.819861.

Xiaoyue Yan, Dazhong Mu, Yang Zhang, and Shiyuan Cao. DAA-Font: Cross-lingual font generation with dual-branch encoding and multi-level axial attention. *IEEE Access*, 14:36673–36687, 2026. doi: 10.1109/ACCESS.2026.3670104.

Zhenhua Yang, Dezhi Peng, Yuxin Kong, Yuyi Zhang, Cong Yao, and Lianwen Jin. FontDiffuser: One-shot font generation via denoising diffusion with multi-scale content aggregation and style contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 6603–6611, 2024. doi: 10.1609/aaai.v38i7.28482.

Jinshan Zeng, Yiyang Yuan, Xijia Wang, Yan Zhang, and Yefei Wang. Cross-lingual font generation via patch-level style contrastive learning and relative position awareness. *Pattern Recognition*, 169: 111937, 2026. doi: 10.1016/j.patcog.2025.111937.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018.

Yufan Zhang, Junkai Man, and Peng Sun. MF-Net: A novel few-shot stylized multilingual font generation method. In *Proceedings of the 30th ACM International Conference on Multimedia*, pp. 2088–2096, 2022. doi: 10.1145/3503161.3548414.

Hui-huang Zhao, Tian-le Ji, Paul L. Rosin, Yu-Kun Lai, Wei-liang Meng, and Yao-nan Wang. Cross-lingual font style transfer with full-domain convolutional attention. *Pattern Recognition*, 155: 110709, 2024. doi: 10.1016/j.patcog.2024.110709.

A Implementation details

Character realization prior. For each episode, candidate families must contain the requested target character and every observed reference character. We exclude the episode family and its explicit weight variants before retrieval. The selected families supply target-character renderings. At each structural module, the frozen content and style encoders separately compute candidate-minus-neutral feature differences. Each branch routes its own differences and predicts an offset; the offsets are combined as

$$\Delta_{\ell,t} = \Delta^{c}_{\ell,t} + \tanh(\zeta_\ell)\Delta^{s}_{\ell,t}. \tag{9}$$

The offset controls a deformable residual update of the skip feature, $\widetilde{U}_{\ell,t} = U_{\ell,t} + P_\ell[\mathrm{DCN}_\ell(U_{\ell,t}, \Delta_{\ell,t}) - U_{\ell,t}]$, where P_ℓ is initialized to zero. Candidate selection is fixed during sampling; routing is recomputed at each denoiser evaluation.

Target-character appearance completion. The local encoder is initialized from the first three blocks of the pretrained style encoder and is trained jointly with the denoiser. For each 96×96 reference, its 24×24 and 12×12 feature maps produce 720 tokens of width 256 after projecting the shallow map. The 12×12 neutral-content map supplies 144 queries, augmented by learned positional vectors. Attention uses four heads of width 64 with head-wise query/key normalization. A residual



Figure 5: Additional qualitative comparisons across Latin, kana, and Bopomofo.

feed-forward block has widths 256, 512, and 256 with a GELU activation. The two output projections produce $M \in \mathbb{R}^{144\times128}$ and $T \in \mathbb{R}^{144\times1024}$. The teacher uses the VGG16 feature map after the second pooling stage, adaptively averaged to 12×12. Its channel means and standard deviations are fixed using 2,048 training targets.

Local conditioning scale. Decoder attention reads global reference tokens and T with shared attention parameters and separate softmax normalizations. For each sample, RMS is computed across token positions and channels as $\mathrm{RMS}(A) = \sqrt{\mathrm{mean}(A^2) + 10^{-6}}$. We set $\beta_0 = 0.8$ and use

$$\kappa_{\ell,t} = \mathrm{sg}\left[\mathrm{clip}\left(\frac{\mathrm{RMS}(A^G_{\ell,t})}{\mathrm{RMS}(A^T_{\ell,t})}, 0.25, 4\right)\right].$$

B Data and evaluation panel

Final dataset contains 137,639 training font–character pairs from 738 fonts after filtering, 7,331 validation pairs from 32 fonts, and 7,088 test pairs from another 32 fonts. The main comparison uses a separately defined panel of 8,638 valid pairs from 47 fonts and 284 target characters: 3,542 Latin, 4,726 kana, and 370 Bopomofo pairs. Glyphs are rendered at 96×96 pixels with a shared font size within each font. Neutral content uses Noto Sans CJK Regular. The accompanying manifest gives the sample IDs and script domains used for the supply testsets panel. The anonymized font inventory omits source and license fields.

The CRP candidate pool contains 423 training font identities. The 47 evaluation font identities are absent from the generator training, validation, and internal test splits and from the candidate pool. All 284 evaluation target characters occur in the generator training data.

The evaluation panel was constructed to include typefaces with pronounced visual characteristics. Its definition and sample IDs are fixed across the supply testsets method comparisons. Public font IDs are pseudonyms; the private mapping to source assets is not part of the anonymous supplement.

C Additional qualitative results

Figure 5 shows additional qualitative comparisons across the target scripts.

D TRAINING, INFERENCE, AND ABLATION PROTOCOLS

The module configurations are Full (HR-Font with CRP and TC), Baseline + CRP (CRP only), Baseline + TC (TC only), and Baseline (neither condition). These configurations are trained as separate models. Ordinary skip connections and global reference conditioning remain active in all four. Evaluation uses the same reference episodes, sampler, and initial diffusion noise, as specified in Table 4. The main HR-Font evaluation uses EMA weights. Training uses eight V100 GPUs and a global batch size of 64. For Full, the base perceptual, completion, and detail loss weights are 0.01, 0.01, and 0.05, respectively, with no offset regularization.

All four models start from the same parent checkpoint with seed 3407 and are trained for 20,000 updates. The pretrained style and content encoders and the VGG teacher remain frozen; the denoiser is updated in every configuration. Disabling TC removes the local reader, while disabling CRP zeros the donor differences and deactivates the structural residual. Consequently, all four objectives include noise-prediction MSE, 0.01-weighted VGG perceptual loss, and 0.05-weighted detail loss; the 0.01-weighted teacher-feature completion loss applies only to Full and Baseline + TC. The completion and detail terms are multiplied by $\min(1, u/1000)$ at update u, and the TC attention contribution is scaled by $0.8\min(1, u/1000)$. No configuration uses offset regularization.

We use AdamW with $(\beta_1, \beta_2) = (0.9, 0.999)$ and $\epsilon = 10^{-8}$. Peak learning rates are 10^{-4} for the local reader and added structural-offset parameters and 2×10^{-5} for the remaining trainable parameters. Weight decay is 0.01 for parameters with more than one dimension and zero otherwise. Learning rates warm up for 500 updates, remain at their peak through update 5,000, decay by cosine to 10% of the peak at update 10,000, and stay there through update 20,000. The eight-GPU setup applies to Full; each separately trained ablation uses four V100 GPUs with two accumulation steps, preserving the global batch size of 64.

Candidate interventions change inference for the trained Full model: Uniform Top-K replaces the final candidate weights with uniform weights, Retrieval Top-1 keeps the highest-ranked admissible candidate, and Random donors draws the same number of candidates from the admissible pool. Reference-budget rows use nested sets of 1, 2, 4, and 8 generation references. Reported Supply testset comparisons use a common 8,638-sample panel; their generation and evaluation reference sets are recorded separately.

Paired font-level comparisons. Each row compares the indicated HR-Font configuration with its comparator. For L1 and LPIPS, the difference is the comparator's score minus the HR-Font score; for SSIM and SC, the order is reversed. Positive values favor the HR-Font configuration in that row. Confidence intervals resample fonts while preserving matched observations for the evaluated trained models. The first three rows use Full with eight generation references; the final row compares HR-Font and FontDiffuser with one reference each. Table 4 defines the configuration names.

Table 6: Paired font-level metric differences. Each reported cell gives an estimate and its 95% confidence interval. Dashes indicate unreported comparisons.

Comparison	ΔL1	ΔSSIM	ΔLPIPS	ΔSC
Full vs. Baseline + CRP	$+0.003235$ $[0.002042, 0.004483]$	$+0.007248$ $[0.004619, 0.010022]$	$+0.015260$ $[0.010619, 0.020987]$	$+0.020736$ $[0.004822, 0.041916]$
Full vs. Baseline + TC	$+0.007014$ $[0.005197, 0.008937]$	$+0.014955$ $[0.010773, 0.019365]$	$+0.024320$ $[0.019679, 0.029137]$	$+0.023277$ $[0.014202, 0.034895]$
Full vs. Baseline	$+0.008237$ $[0.006180, 0.010370]$	$+0.017958$ $[0.013223, 0.023086]$	$+0.033164$ $[0.025182, 0.041909]$	$+0.042736$ $[0.022889, 0.065857]$
HR-Font vs. FontDiffuser (1 ref. each)	–	–	–	$+0.049523$ $[0.023096, 0.082772]$

E METRICS AND EXPERT RANKING

L1, SSIM, and AlexNet-based LPIPS compare generated glyphs with paired targets. Style Compatibility (SC) is the raw cosine similarity between a generated glyph embedding and the nor-

malized prototype of its Chinese evaluation references. Evaluation references are fixed across methods and generation budgets. With unit-normalized encoder ϕ, the score is $\mathrm{SC}(\hat{y}, R) = \langle E_{\mathrm{SC}}(\hat{y}), \mathrm{normalize}(\sum_{r \in R} E_{\mathrm{SC}}(r)) \rangle$. The encoder is a ResNet-18 with its classifier removed and its 512-dimensional pooled feature ℓ_2-normalized. SC uses native 96×96 RGB inputs scaled to $[0, 1]$ and normalized with ImageNet channel statistics. The eight evaluation references are fixed across methods and generation reference budgets. The encoder is trained independently, frozen for evaluation, and used after generator checkpoint selection. The supplement provides the encoder definition, training loss, and scoring code; the trained encoder weights are not included. Reported metrics are averaged within each font and then equally across fonts. DINOv2 distance is an additional diagnostic in the Supply testsets tables below; lower values indicate greater similarity to the paired target.

F Expert ranking protocol

Twenty-eight designers and domain experts each completed 30 formal Latin-character items. A separate expert group selected the items from ground-truth glyphs before inspecting generated outputs; selectors and raters did not overlap, and neither group included authors or model developers. Each item showed Chinese style references, five generated candidates, and the ground-truth glyph. Method identities and the ground-truth label were concealed, and candidate positions were randomized. Participants ranked all six candidates for consistency of stroke weight, proportions, contour treatment, and decoration with the references. Practice items and two style-mismatch checks were separate from the 30 formal items. A participant was excluded only if the mismatched candidate was placed in the top two on both checks; no participant met that rule. Repeated responses for the same participant, item, and candidate were resolved by keeping the final response. The 840 complete rankings are the unit for Top-2 inclusion rate; confidence intervals resample items, retaining their candidate groups. The association analysis uses 30 items times five generated methods, with item-level bootstrap resampling. The supply testsets DINOv2 distance association with expert mean rank is $\tau_b = 0.281$ with a 95% item-bootstrap interval of $[0.199, 0.370]$; this is an additional diagnostic and does not replace the main paper metrics.

Table 7 reports within-item ordering agreement for pairs of generated candidates.

Table 7: Within-item ordering agreement with expert rankings (%). The second interval resamples items; the third clusters by font.

Metric	Agreement [95% CI]	Font-clustered 95% CI
L1	43.2 [36.1, 50.3]	[36.3, 49.2]
SSIM	47.9 [40.0, 55.9]	[39.6, 55.6]
LPIPS	74.0 [67.3, 80.3]	[66.3, 81.1]
SC	66.8 [60.4, 73.0]	[60.5, 73.2]

G Additional results

Table 8 reports overall results on the same 8,638-pair panel used in the main comparisons. Metrics are averaged over characters within each font and then across fonts with equal weight.

Tables 9, 10, and 11 report font-macro averages within each script.

Table 8: Overall results on the 8,638-pair test panel. Values are font-macro averages.

Method	SC ↑	LPIPS ↓	L1 ↓	SSIM ↑	DINOv2 dist. ↓
HR-Font (8 refs.)	0.699777	0.176435	0.084549	0.739638	0.203033
Baseline + CRP	0.679041	0.191694	0.087784	0.732390	0.222579
Baseline + TC	0.676501	0.200755	0.091563	0.724683	0.224656
Baseline	0.657041	0.209599	0.092786	0.721680	0.235707
FCA-GAN	0.475116	0.228859	0.063130	0.783203	0.281371
FTransGAN	0.480683	0.229734	0.063586	0.782643	0.270595
GAR-Font	0.584269	0.205034	0.071588	0.767646	0.201781
FontDiffuser	0.613962	0.186309	0.082526	0.741735	0.196962
HR-Font (1 ref.)	0.663486	0.191209	0.087498	0.733956	0.217314
HR-Font (2 refs.)	0.688275	0.180277	0.084951	0.739869	0.207020
HR-Font (4 refs.)	0.696130	0.179004	0.084269	0.740546	0.206400
Retrieval Top-1	0.698470	0.171600	0.080835	0.748479	0.202070
Uniform Top-K	0.699714	0.176462	0.084568	0.739591	0.203254
Random donors	0.695751	0.180701	0.087100	0.733961	0.204132

Table 9: Font-macro results for Latin targets. A dash indicates an unreported value.

Method	SC ↑	LPIPS ↓	L1 ↓	SSIM ↑	DINOv2 dist. ↓
HR-Font (8 refs.)	0.691520	0.160517	0.069967	0.782342	0.207393
Baseline + CRP	0.671761	0.178783	0.073333	0.774418	0.229594
Baseline + TC	0.666837	0.186816	0.077313	0.766203	0.230997
Baseline	0.647615	0.196589	0.078455	0.763309	0.242739
FCA-GAN	–	–	–	–	–
FTransGAN	–	–	–	–	–
GAR-Font	–	–	–	–	–
FontDiffuser	–	–	–	–	–
HR-Font (1 ref.)	0.659572	0.175662	0.073493	0.774531	0.221898
HR-Font (2 refs.)	0.681828	0.164034	0.070478	0.782074	0.211431
HR-Font (4 refs.)	0.687552	0.163285	0.069945	0.782801	0.211056
Retrieval Top-1	0.689803	0.156572	0.067247	0.788702	0.206411
Uniform Top-K	0.691406	0.160524	0.069982	0.782312	0.207707
Random donors	0.688611	0.164462	0.072293	0.777000	0.208255

Table 10: Font-macro results for kana targets. A dash indicates an unreported value.

Method	SC ↑	LPIPS ↓	L1 ↓	SSIM ↑	DINOv2 dist. ↓
HR-Font (8 refs.)	0.728230	0.187637	0.101035	0.687152	0.188929
Baseline + CRP	0.704637	0.203021	0.105269	0.677979	0.206799
Baseline + TC	0.703075	0.213224	0.109072	0.669586	0.211350
Baseline	0.683246	0.219467	0.110369	0.666546	0.219083
FCA-GAN	–	–	–	–	–
FTransGAN	–	–	–	–	–
GAR-Font	–	–	–	–	–
FontDiffuser	–	–	–	–	–
HR-Font (1 ref.)	0.695805	0.199266	0.103793	0.682446	0.203507
HR-Font (2 refs.)	0.721054	0.191327	0.101872	0.686639	0.194018
HR-Font (4 refs.)	0.726365	0.188603	0.100599	0.688178	0.191346
Retrieval Top-1	0.727338	0.179995	0.095962	0.699564	0.187176
Uniform Top-K	0.728209	0.187702	0.101065	0.687075	0.188891
Random donors	0.721467	0.192157	0.103706	0.681136	0.190172

Table 11: Font-macro results for Bopomofo targets. A dash indicates an unreported value.

Method	SC ↑	LPIPS ↓	L1 ↓	SSIM ↑	DINOv2 dist. ↓
HR-Font (8 refs.)	0.708496	0.186874	0.097140	0.706687	0.185488
Baseline + CRP	0.705917	0.199103	0.102328	0.697266	0.202439
Baseline + TC	0.684516	0.220027	0.107800	0.686193	0.220951
Baseline	0.662717	0.220961	0.109398	0.682988	0.226036
FCA-GAN	–	–	–	–	–
FTransGAN	–	–	–	–	–
GAR-Font	–	–	–	–	–
FontDiffuser	–	–	–	–	–
HR-Font (1 ref.)	0.670410	0.203498	0.102929	0.697207	0.205359
HR-Font (2 refs.)	0.705616	0.193140	0.099096	0.703404	0.192042
HR-Font (4 refs.)	0.712734	0.188296	0.097165	0.706408	0.186764
Retrieval Top-1	0.707853	0.183562	0.093846	0.714576	0.186183
Uniform Top-K	0.708705	0.186934	0.097136	0.706672	0.185280
Random donors	0.699026	0.190602	0.099158	0.702643	0.188963