

# QwenStyle: Content-Preserving Style Transfer with Qwen-Image-Edit

Shiwen Zhang

Haibin Huang

Chi Zhang

Xuelong Li

Institute of Artificial Intelligence (TeleAI), China Telecom



Figure 1. QwenStyle accepts style and content references for content-preserving style transfer, while maintaining high aesthetics merit. QwenStyle is the first content-preserving style transfer model built on Qwen-Image.

## Abstract

*Content-Preserving Style transfer, given content and style references, remains challenging for Diffusion Transformers (DiTs) due to its internal entangled content and style features. In this technical report, we propose the first content-preserving style transfer model trained on Qwen-Image-Edit, which activates Qwen-Image-Edit's strong content preservation and style customization capability. We collected and filtered high quality data of limited specific styles and synthesized triplets with thousands categories of style images in-the-wild. We introduce the Curriculum Continual Learning framework to train QwenStyle with such mixture of clean and noisy triplets, which enables QwenStyle to generalize to unseen styles without degradation of the precise content preservation capability. Our QwenStyle V1 achieves state-of-the-art performance in three core metrics: style similarity, content consistency, and aesthetic quality.*

## 1. Introduction

In this technical report<sup>1</sup>, for the first time, we introduce content-preserving style transfer functionality to Qwen-Image-Edit model [23]. Unlike SDXL [13] based style transfer models [21, 27], where content and style could be separated by UNet Disentangle Law [25, 26], we could not find such effective disentanglements in Diffusion Transformers [2, 6, 12]. Thus it becomes a challenging task for content-preserving style transfer on DiT models, which requires the model to transfer style cues from style reference to content reference while preserving the characteristics of the content reference. Qwen-Image and Qwen-Image-Edit series [23] are the state-of-the-art open-sourced text-to-image and image editing models, built upon powerful Qwen2.5-VL [1] with strong language understanding

<sup>1</sup>The codes and models are released at <https://github.com/witcherofresearch/Qwen-Image-Style-Transfer>. If you have any questions regarding this paper, feel free to contact me via my personal email [witcherofresearch@gmail.com](mailto:witcherofresearch@gmail.com)

and MMDiT [2] for image generation. Yet they don’t support style transfer with content and style references at this moment.

Since Qwen-Image-Edit has already applied MS-RoPE, an effective variant of RoPE [18], to distinguish different reference images, we believe the core solutions to content-preserving style transfer on Qwen-Image-Edit are curated training triplets and multi-stage curriculum learning paradigm. Our QwenStyle achieves state-of-the-art performance in terms of style similarity, content preservation and aesthetic merit. We demonstrate a few examples in Figure 1, where our QwenStyle could transfer various styles while preserving complex characteristics.

## 2. QwenStyle

### 2.1. Triplet Training Dataset Construction

Unlike subject-driven image pair/triplet data [19, 28], which naturally exists in videos or photo albums, style triplets are rare in real world. We collected [Style Ref, Content Ref, Target] image triplets from a dataset [17] sampled from GPT-4O [5] and some Loras from open-source community, and purified them with data cleaning. However, such collection is expensive and we only obtain 30 style categories. The model trained on such limited style categories turns out to generalize poorly to unseen styles. Thus we introduce a reverse triplet synthetic framework inspired by [22] to generate training triplets from style images in-the-wild [7], where different style images are organized into noisy style clusters. The synthesis framework is shown in Figure 2. Given a stylized target image, we convert it to a photographic content reference with our previously trained image editing model based on FLUX-dev [6] (due to historical reasons and limited time, we did not train such a realism converter with Qwen-Image, which we will improve in the future). In addition, we generate a style reference image with CDST [27], which uses Dino v2 [11] to extract style representations. We sample a random prompt from a generated prompt bank, where no human appears in the prompts because we found that SDXL-based CDST suffers from leakage problem which may pollute human identity in the generated image. For simplicity, we call the collected clean triplet dataset  $D_{collected}$  and the synthetic dataset  $D_{synthetic}$  in the following sections. By matching every pair of stylized images in the same style cluster (one is style reference, the other one is the target image), we obtain 300k triplets in  $D_{collected}$ . With a similar matching strategy, we acquire 1 million triplets in  $D_{synthetic}$ .

### 2.2. Content-Preserving Style Transfer with Curriculum Continual Learning

Due to different qualities of data sources and different learning difficulties of various style categories, we found

it obliged to train QwenStyle in a multi-stage paradigm. Specifically, we introduce a Curriculum Continual Learning paradigm to tackle this challenging task.

1. In the first stage, we train QwenStyle with  $D_{collected}$ , which is denoted as  $D_1$ , to activate the content-preserving style transfer capability of Qwen-Image-Edit. However, we found the trained model  $Q_1$  has two problems:
  - **problem 1.** The model could not preserve subtle characteristics in the content reference image, for example, multiple facial identities.
  - **problem 2.** The model’s generalization capability is poor and could not effectively adapt to out-of-distribution style references.
2. In the second stage, we aim to improve characteristics preservation of  $Q_1$ . We found that even in triplets from  $D_{collected}$ , there are still many inconsistent content reference which alters the characteristics of target image. We manually classified the triplets into high and low content similarity. We increase the ratio of high content similarity data and denote the new data  $D_2$ . Initialized with model  $Q_1$  from stage 1, we finetune Qwen-Image-Edit with  $D_2$  and obtain model  $Q_2$ . Results on the validation set demonstrate that  $Q_2$  could preserves the characteristics of content reference in a very precise manner, much better than  $Q_1$ . Thus, problem 1 is solved.
3. In the third stage, we focus on improving style consistency while keeping the content preserving capability. Though we could match all the possible triplets of  $D_{synthetic}$  to obtain 1 million data, the quantity does not help the training and leads to degradation of content consistency instead. The content reference from  $D_{synthetic}$ , though overall roughly preserves the characteristics of target, usually alters some details, for example, some small objects or the facial identities in the target image etc. Thus we need to keep a low ratio of  $D_{synthetic}$  during training. In order to avoid catastrophic forgetting on precise content preservation capability of model  $Q_2$  from stage 2, we need to preserve certain amount of training data  $D_2$ , which is merged with a low ratio of  $D_{synthetic}$  to form training data for Stage 3, denoted  $D_3$ . Initialized with  $Q_2$ , we introduce the Curriculum Continual Learning paradigm to train the final model  $Q_3$ . Through experiments on validation set, we found that for  $Q_3$ , the degradation of content consistency is acceptable and generalization ability on out-of-distribution styles is significantly improved. Thus, problem 1 is almost solved and problem 2 is alleviated somewhat.

### 2.3. Training and Inference

We found that training all the parameters of Qwen-Image MMDiT [23] does not lead to significant performance advantage over training a Lora [4] with rectified flow-

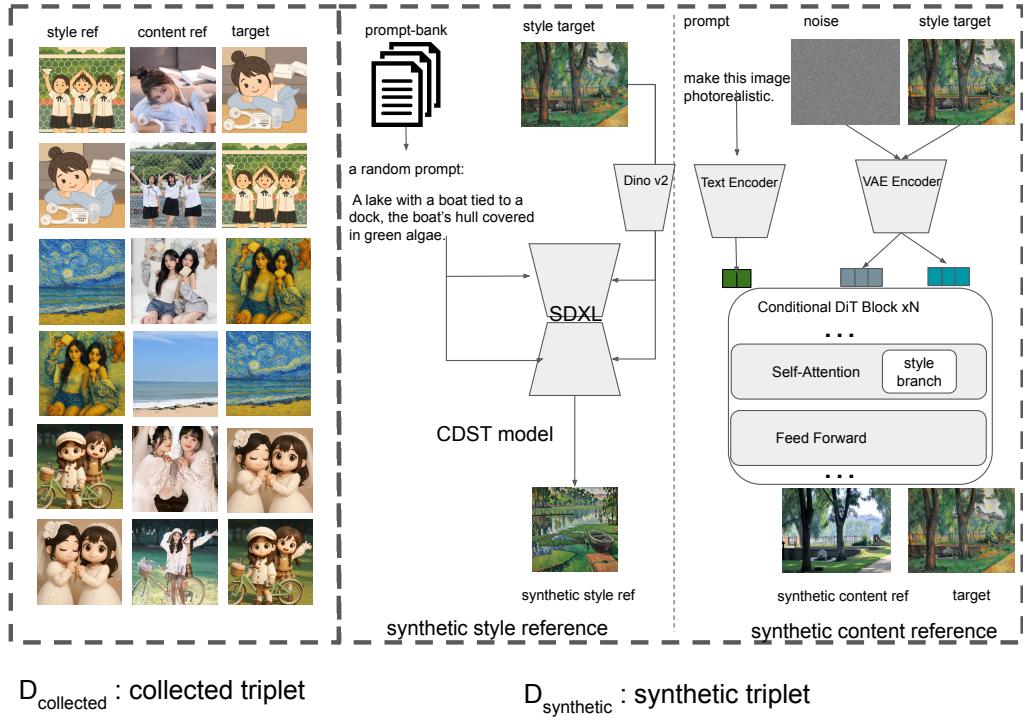


Figure 2. Collected triplets  $D_{collected}$  and synthetic triplets  $D_{synthetic}$

matching [8, 9], thus we open-source the Lora version of QwenStyle. During training, the prompt is always

- "Style Transfer the style of Figure 2 to Figure 1, and keep the content and characteristics of Figure 1."

We recommend using the same prompt during inference, though QwenStyle V1 also supports other prompts, for example,

- "Transfer Figure 1 into XXX material."
- "Transfer Figure 1 into XXX style."

For example, if the style/material the style reference is known, we could use the prompt "Transfer Figure 1 into metal material." or 'Transfer Figure 1 into Van-Gogh style' to strengthen the style fidelity of style reference if the default prompt does not lead to successful edits, though none of the generated images in this technical report use such prompt formats. Please check our online free demo to explore more examples.

In addition, we found that during inference, it is crucial to keep the same height and width for the content reference and the generated image. For style reference, we simply resize it into a square whose height and width are  $\min(H_{target}, W_{target})$ .

### 3. Experiments

**Implementation Details.** We adopt Lora [4] to Qwen-Image-Edit-2509 [23] to train QwenStyle V1. The ranks for Lora is 32. Gradient Checkpointing [3] is applied to

---

#### Algorithm 1 Content-Preserving Style Transfer with Curriculum Continual Learning

---

- 1: **Input:**
- 2:  $D_1$ : triplets of natural distribution  $D_{collected}$ ,
- 3:  $D_2$ : increase the ratio of triplets with high content consistency in  $D_1$ ,
- 4:  $D_3$ :  $D_2$  mixed with a low ratio of  $D_{synthetic}$ .
- 5:
- 6: **Output:** QwenStyle V1
- 7:
- 8: **Procedure:**
- 9: Train Lora  $Q_1 \leftarrow \{D_1, \text{Qwen-Image-Edit}\}$
- 10: Train Lora  $Q_2 \leftarrow \{D_2, Q_1\}$
- 11: Train Lora  $Q_3 \leftarrow \{D_3, Q_2\}$
- 12: QwenStyle V1  $\leftarrow Q_3$
- 13: **return** QwenStyle V1

---

save memory and the model is trained with min-edge=1024. Our model is trained with 4 H100 GPUs, batch size is 1 for each GPU, learning rate is 1e-4.

**Evaluation Benchmark.** We select 50 style references and 40 content references of different ratios, mutually pair each of them to generate 2000 style-content pairs for testing. We further select 10 style references and 10 content references as validation set. The style references cover diverse

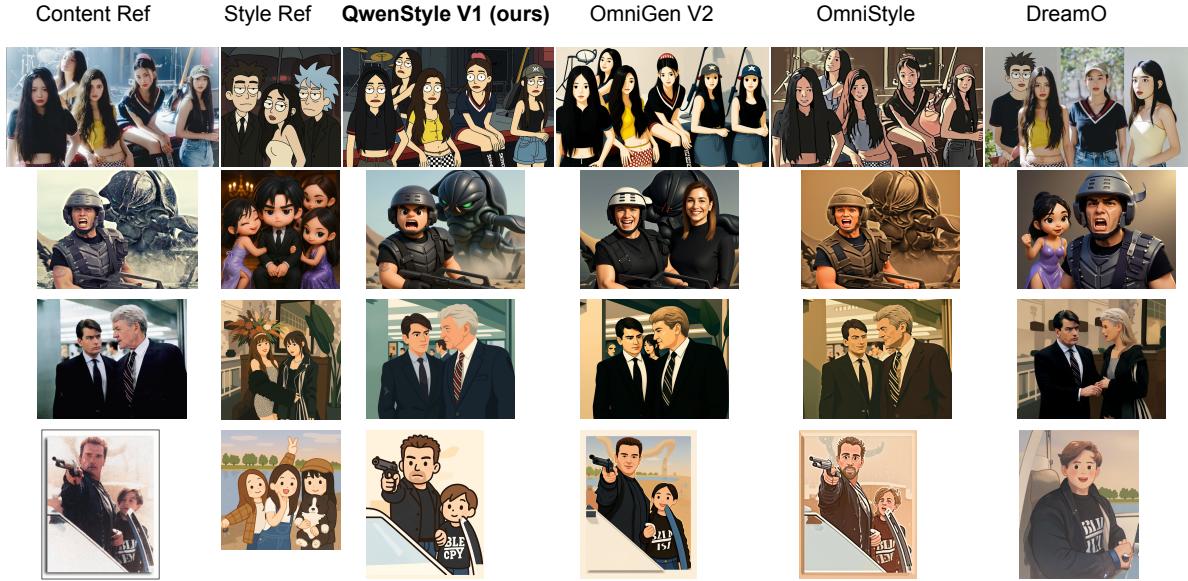


Figure 3. Qualitative Comparison with State-of-the-art Style Transfer Models.

Model	Style Similarity CSD Score $\uparrow$	Content Preservation CPC Score@0.5 $\uparrow$	Content Preservation CPC Score@0.3:0.9 $\uparrow$	Aesthetic Score $\uparrow$
OmniStyle [20]	0.447	0.194	0.163	5.881
OmniGen-v2 [24]	0.462	0.243	0.166	5.843
DreamO [10]	0.402	0.193	0.102	6.149
<b>Style-CCL (ours)</b>	<b>0.577</b>	<b>0.441</b>	<b>0.304</b>	<b>6.317</b>

Table 1. Quantitative comparison of our Style-CCL with previous state-of-the-art style transfer methods. The best score is stressed by bold font and the second best score is marked by underline.

style genres and the content references include different number of persons with diverse gestures, scenes/buildings and subjects in complex scenarios.

**Evaluation Metrics.** We evaluate our method with the following metrics. For **Style Consistency**, we use CSD Score [16] to measure the style similarity between the style reference and the generated image. For **Aesthetics**, we use the LAION Aesthetics Predictor [15] to estimate the aesthetic quality of the generated image. For **Content Preservation**, we propose a new *Content Preservation Cut-Off Score* (CPC Score), which augments the original Content Preservation Score [20] with a style consistency threshold. Intuitively, a model that simply replicates the content reference without transferring style would receive an artificially high content score. To avoid this, we first use Qwen-VL [1] to generate a detailed caption  $T_{vlm}$  for the content reference image  $I_{content}$ , and compute the CLIP score [14] between  $T_{vlm}$  and the generated image  $I_{res}$ . We then compute the CSD Score between  $I_{style}$  and  $I_{res}$ ; if this score falls below a threshold, the CLIP score is set to zero as a penalty.

$$CPC@thresh = \begin{cases} CLIP(I_{res}, T_{vlm}), & \text{if } CSD(I_{res}, I_{style}) \geq \text{thresh} \\ 0, & \text{if } CSD(I_{res}, I_{style}) < \text{thresh} \end{cases} \quad (1)$$

### 3.1. Comparison with State-of-the-art Methods

**Quantitative Comparison** We quantitatively compare our QwenStyle V1 with multiple DiT-based style transfer models in Table 1, from the aspects of style similarity, content preservation and aesthetics score. Our QwenStyle V1 demonstrates significant advantages over previous models.

**Qualitative Comparison** We present qualitative visual comparison with some representative DiT-based style transfer models in Figure 3, where our QwenStyle V1 demonstrates superior style similarity and content consistency than previous models, while maintaining high aesthetics values.

**Limitations.** QwenStyle V1 is still an experimental attempt to tackle the challenging content-preserving style transfer task. There is still considerable potential to improve its generalization capability on out-of-distribution styles.

## 4. Conclusion

We present QwenStyle V1, the first content-preserving style transfer model trained on Qwen-Image-Edit. With a Three-Stage Curriculum Continual Learning framework introduced, QwenStyle could preserve content characteristics and generalize to unseen styles. QwenStyle V1 achieves new state-of-the-art performance on style similarity, content preservation and aesthetics score.

## References

- [1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. [1](#), [4](#)
- [2] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. [1](#), [2](#)
- [3] Andreas Griewank and Andrea Walther. Algorithm 799: revolve: an implementation of checkpointing for the reverse or adjoint mode of computational differentiation. *ACM Transactions on Mathematical Software (TOMS)*, 26(1):19–45, 2000. [3](#)
- [4] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. [2](#), [3](#)
- [5] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. [2](#)
- [6] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. [1](#), [2](#)
- [7] Wen Li, Muyuan Fang, Cheng Zou, Biao Gong, Ruobing Zheng, Meng Wang, Jingdong Chen, and Ming Yang. Style-tokenizer: Defining image style by a single instance for controlling diffusion models. In *European Conference on Computer Vision*, pages 110–126. Springer, 2024. [2](#)
- [8] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. [3](#)
- [9] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. [3](#)
- [10] Chong Mou, Yanze Wu, Wenxu Wu, Zinan Guo, Pengze Zhang, Yufeng Cheng, Yiming Luo, Fei Ding, Shiwen Zhang, Xinghui Li, et al. Dreamo: A unified framework for image customization. *arXiv preprint arXiv:2504.16915*, 2025. [4](#)
- [11] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [2](#)
- [12] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. [1](#)
- [13] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. [1](#)
- [14] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021. [4](#)
- [15] Christoph Schuhmann and Romain Beaumont. Laion-aesthetics. *LAION AI*, 2022. [4](#)
- [16] Gowthami Somepalli, Anubhav Gupta, Kamal Gupta, Shramay Palta, Micah Goldblum, Jonas Geiping, Abhinav Shrivastava, and Tom Goldstein. Measuring style similarity in diffusion models. *arXiv preprint arXiv:2404.01292*, 2024. [4](#)
- [17] Yiren Song, Cheng Liu, and Mike Zheng Shou. Omniconsistency: Learning style-agnostic consistency from paired stylization data. *arXiv preprint arXiv:2505.18445*, 2025. [2](#)
- [18] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. [2](#)
- [19] Zhenxiong Tan, Songhua Liu, Xingyi Yang, Qiaochu Xue, and Xinchao Wang. Ominicontrol: Minimal and universal control for diffusion transformer. *arXiv preprint arXiv:2411.15098*, 3, 2024. [2](#)
- [20] Ye Wang, Ruiqi Liu, Jiang Lin, Fei Liu, Zili Yi, Yilin Wang, and Rui Ma. Omnistyle: Filtering high quality style transfer data at scale. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7847–7856, 2025. [4](#)
- [21] Zhouxia Wang, Xintao Wang, Liangbin Xie, Zhongang Qi, Ying Shan, Wenping Wang, and Ping Luo. Styleadapter: A unified stylized image generation model. *arXiv preprint arXiv:2309.01770*, 2023. [1](#)
- [22] Zhizhong Wang, Lei Zhao, and Wei Xing. Stylediffusion: Controllable disentangled style transfer via diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7677–7689, 2023. [2](#)
- [23] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. [1](#), [2](#), [3](#)
- [24] Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan Jiang, Yexin Liu, Junjie Zhou, et al. Omnipgen2: Exploration to advanced multimodal generation. *arXiv preprint arXiv:2506.18871*, 2025. [4](#)
- [25] Shiwen Zhang. Fast Imagic: Solving Overfitting in Text-guided Image Editing via Disentangled UNet with Forgetting Mechanism and Unified Vision-Language Optimization. In *PMLR*, 2024. [1](#)

- [26] Shiwen Zhang, Shuai Xiao, and Weilin Huang. Forgedit: Text guided image editing via learning and forgetting. *arXiv preprint arXiv:2309.10556*, 2023. [1](#)
- [27] Shiwen Zhang, Zhuowei Chen, Lang Chen, and Yanze Wu. Cdst: Color disentangled style transfer for universal style reference customization. *arXiv preprint arXiv:2506.13770*, 2025. [1](#), [2](#)
- [28] Yuxuan Zhang, Yirui Yuan, Yiren Song, Haofan Wang, and Jiaming Liu. Easycontrol: Adding efficient and flexible control for diffusion transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19513–19524, 2025. [2](#)