Multi-Attribute Guided Painting Generation

Minxuan Lin NLPR, CASIA & School of AI, UCAS linminxuan2018@ia.ac.cn Yingying Deng School of AI, UCAS & NLPR, CASIA dengyingying2017@ia.ac.cn

Fan Tang Fosafer tangfan@fosafer.com Weiming Dong * NLPR, CASIA & CASIA-LLVision Joint Lab weiming.dong@ia.ac.cn Changsheng Xu NLPR, CASIA csxu@nlpr.ia.ac.cn

Abstract

Controllable painting generation plays a pivotal role in image stylization. Currently, the control way of style transfer is subject to exemplar-based reference or a random one-hot vector guidance. Few works focus on decoupling the intrinsic properties of painting as control conditions, e.g., artist, genre and period. Under this circumstance, we propose a novel framework adopting multiple attributes from the painting to control the stylized results. An asymmetrical cycle structure is equipped to preserve the fidelity, associating with style preserving and attribute regression loss to keep the unique distinction of colors and textures between domains. Several qualitative and quantitative results demonstrate the effect of the combinations of multiple attributes and achieve satisfactory performance.

1 Introduction

Painting generation is a technique used to create art by synthesizing style patterns from given painting image(s) evenly over a natural image while maintaining its original content structure. For art creation, every painter has a unique interpretation of painting style. Moreover, the painting preferences (e.g., the usage of strokes and colors) change at different periods. Commonly, the attributes like *artist*, *genre* and *period* are regarded as crucial factors to represent paintings, which will also affect the visual appearance of style transfer results. Figure 1 shows the art stream of individual painting preference along with genres and career periods. The disparities of colors and brush strokes reflect in different genres and periods, forming some representative features which belong to specific attribute combinations. Thus, using these attributes as control conditions is an



Figure 1. The artwork stream of diverse artists. The drawing habits of color and brush stroke gradually alter over period and genre.

acceptable measure to guide painting generation. However, most existing methods [6, 7] mainly focus on increasing the quality of stylized results but are lack of high controllability in appearance guided by these properties.

To address this problem, we propose a practical approach, *multi-attribute guided painting generation*, for artistic stylized image generation where painting appearance can be easily controlled by user-assigned attributes. Specifically, motivated by CycleGAN [10], an *asymmetrical* cycle structure is adopted as content branch to generate stylized results. With regard to the multi-attribute as style control conditions, rather than simply stacking the conditions into the feature layers, a congregate multi-attribute vector is parsed by a multi-layer preception network as style guidance to produce AdaIN [4] parameters. To enhance the distinction, we utilize a multi-task discriminator and style preserving constraint to enlarge the color and texture gaps among different domains.

^{*}Corresponding author



Figure 2. The asymmetric network of our proposed method. (a) A forward flow conditioned with attributes. (b) A backward synthesis stream. The legend is listed on the top right corner.

2 Related works

Painting generation Since Gatys et al. [3] proposed convolution based style transfer method, painting generation has been widely investigated. Johnson et al. [5] adopted perception constraint and achieved the real-time synthesis of single painting. Sanakoyeu et al. [7] accomplished multi-style generation by using a style-aware adversarial network. Several works [4, 6, 2] emphasized the feasibility of arbitrary style transfer. The capacity of maintaining colors and textures from random style images was still reserved without retraining. However, the controllability of those methods is limited by exemplar-guidance based generation mechanism. Zhu et al. [10] utilized the cycle consistency to stylize unpaired images with the average domain styles. Choi et al. [1] presented a unified generator conditioned with attribute to exchange styles among multiple domains. However, homogeneous attributes are lack of the maneuverability from various aspects.

Painting attributes parsing The analysis of intrinsic attributes of painting is challenging, some inherent properties (e.g., artist, genre and period) of artworks are extracted for deep exploration. Though the creation techniques change throughout painters' careers, the common style trait is still retained. Van Noord et al. [8] adopted style features from different sub-regions of painting to recognize artists. Taking the period factor into consideration, Yang et al. [9] discovered that paintings' distribution organized according to style. These attributes carry enough style information.

3 Framework

Assuming $x \in \mathcal{X}$ is an image from *content domain* and the attribute set c = (a, p, g) represents three natural properties of paintings: Artist, Period and Genre. The goal of model is to generate the stylized images \tilde{y} , carrying the corresponding characteristics for specific attribute set c. To this end, as shown in Figure 2, we propose an asymmetric cycle synthesis framework composed of two unidirectional generators G and F to make controllable painting generation. Meanwhile, two inconsistent discriminators D_y and D_x take responsibility to distinguish the real and fake images and regress the style attributes.

Conditional forward synthesis Figure 2(a) depicts the forward generation flow. The content image x is fed into the multiple convolution blocks for downsampling, then several residual blocks are used to encode the content feature maps. The attributes of artist a and period p are expressed as one-hot labels. We specially design the representation of genre g, a small perturbation sampled from a gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ and add it to the non-zero dimension of one-hot vector as the expression of genre. This operation aims to introduce random variations to improve the robustness of model for constructing a smoother embedding space.

After cascading these one-hot labels which forms the attribute aggregation c, we use a MLP network to parse the condition and unfolds c to a high dimension space as AdaIN [4] parameters. The adaptive instance norm (AdaIN) technique is designed for aligning the second-order statistics of content and reference images which reflect the style. Thus, some residual blocks are normalized by AdaIN to produce the stylized content features. Finally, instead of deconvolution, we use upsampling and convolution operations to decode the features to generate the stylized image \tilde{y} .

Based on the adversarial structure, the generated image \tilde{y} and the real style image y are alternately sent into a multitask discriminator D_y . For the fidelity of stylized image, the *forward adversarial loss* is used to discriminate the source



Figure 3. Qualitative comparison. (a) The input images. (b)-(e) Stylized results from different models (from left to right): StarGAN [1], CycleGAN [10] and ours.

of input and guide the forward generator G:

$$\mathcal{L}_{adv-f} = \mathbb{E}_y[\log D_y(y)] + \mathbb{E}_{x,c}[\log(1 - D_y(G(x,c)))],$$
(1)

where the task of D_y is the regression of attribute *c*. Through estimating the conditions, the generator *G* can better understand the similarities and differences between styles. The *attribute regression loss* is formulated as:

$$\mathcal{L}_{reg} = \mathbb{E}_y[\mathcal{L}_{D_y}(c|y)] + \mathbb{E}_{x,c}[\mathcal{L}_{D_y}(c|G(x,c))], \quad (2)$$

where \mathcal{L}_{D_y} is the cross-entropy loss.

To further ensure the style consistency of synthetic images and real paintings, we employ the *style preserving loss* to measure the similarity. The feature maps of a frozen parameter VGG16 network are used to calculate the L1distance of gram matrix between \tilde{y} and y:

$$\mathcal{L}_{sp} = \sum_{i=1}^{4} \| \mathrm{VGG}_{feat}^i(\tilde{y}) - \mathrm{VGG}_{feat}^i(y) \|_1.$$
(3)

Asymmetric backward generation Unlike the structure in forward flow, without condition, the backward generator F pulls the attributes away from the \tilde{y} to recover the original content image x. In Figure 2(b), a decoder-encoder structure constitutes F. We obtain the reconstructed image \hat{x} by using reverse recovery. The content of original image and the corresponding reconstruction are aligned in pixel level by the L1 metric. The full *reconstruction loss* is formulated as:

$$\mathcal{L}_{rec} = \|F(G(x,c)) - x\|_1 + \|G(F(y),c) - y\|_1 + \|G(y,c) - y\|_1 + \|F(x) - x\|_1.$$
(4)

Moreover, the *backward adversarial loss* is also applied as well as the forward one. Without attribute regression constraint, the discriminator D_x only estimates the source of images:

$$\mathcal{L}_{adv_b} = \mathbb{E}_x[\log D_x(x)] \\ + \mathbb{E}_y[\log(1 - D_x(F(y)))].$$
(5)



Figure 4. Painting generation of different genres. The results from Picasso or Van gogh are stylized with three genre attributes.

Full loss We combine all loss terms as the final objective:

$$\min_{G,F} \max_{D_y,D_x} \mathcal{L}_{full} = \mathcal{L}_{adv_f} + \mathcal{L}_{adv_b} + \lambda_{rec} \mathcal{L}_{rec} + \lambda_{reg} \mathcal{L}_{reg} + \lambda_s \mathcal{L}_{sp},$$
(6)

where we set $\lambda_{rec} = 10, \lambda_{reg} = 1, \lambda_s = 1e^{-4}$ in training.

4 Experiments

Experimental setup We pick the artworks of Picasso, Cézanne, Monet and Van Gogh as traning data to test the algorithm. The artworks and corresponding painting attributes are downloaded from *Wikiart.com*. For each artist, we select two representative periods which appear obvious changes in the art history, e.g., *Cubist* and *Neoclassicist* period of Picasso, *Mature* and *Final* period of Cézanne. If an artist's works are lack of clear period definition, we just reserve paintings in the early and late time of his/her career. For genres, *Impressionism, Cubism* and *Surrealism* are used. Content images are collected from *Pexels.com*. The resolution of our input and output images is 256×256 .



(a) The Monet's impressionism results.

(b) The Picasso's surrealism results.

Figure 5. Painting generation of different periods. The stylized images reveal the disparity between early and late periods of an artist.

Method	Photo2Monet		Photo2Picasso		Photo2Vangogh		Photo2Cézanne	
	IS	Acc	IS	Acc	IS	Acc	IS	Acc
StarGAN	1.53	0.45	1.37	0.07	1.17	0.87	1.59	0.12
Ours	1.30	0.90	1.87	0.62	1.52	0.73	2.23	0.55



Qualitative evaluation In Figure 3, the quality of images stylized by StarGAN [1], CycleGAN [10] and ours are compared under the control of artist attributes. We retrain these models on our dataset. The period and genre attributes of our model are fixed for a fair comparison. We can see that StarGAN easily confuses the styles between domains so that the results are affected by other artists' style. Due to the adoption of deconvolution, CycleGAN produces more artifacts. Our method successfully perserves the variations among domains and obtains high quality results.

Figure 4 demonstrates the controllability of our method for genres. Conditioned by artist and period, the textures and colors of stylized images change with genres. We use fixed one-hot vector for genre in test. The main characteristics of genre are extracted (e.g., the green fits with the habit of Cézanne in impressionism) and are expressed well. Particularly, we also create some zero-shot results (e.g., Cézanne's Surrealism) by mixing the attributes. Figure 5 displays the influence of period. Considering the brush stroke strongly associating with the unaltered genre, colors become the major variations. For instance, the Picasso's surrealism results in late period are more bright and abstract than the early ones.

Quantitative evaluation We calculate the artist classification accuracy and the IS metric for StarGAN and our method, which both contain a unified generator for a fair comparison. The accuracy computed by a finetuned ResNet-18 network examines the similarity between stylized images and artist domains. The IS metric indicates the reality of stylized images, testing on a finetuned Inception-V3 network. As shown in Table 1, our method can outperform the StarGAN method [1].

5 Conclusions

We propose a multi-attribute guided stylization method to increase more controllability for painting generation. Internal painting properties (e.g., artist, genre and period) are utilized as conditions. An asymmetrical cycle structure with control branch constitutes our gframework. The qualitative and quantitative evaluations show the superiority of our model over the state-of-the-art methods.

Acknowledgements

This work was supported by National Natural Science Foundation of China under nos. 61832016 and 61672520.

References

- Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. StarGAN: Unified generative adversarial networks for multidomain image-to-image translation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8789–8797, 2018.
- [2] L. Doyle, F. Anderson, E. Choy, and D. Mould. Automated pebble mosaic stylization of images. *Computational Visual Media*, 5(1):33–44, 2019.
- [3] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), pages 2414–2423, 2016.
- [4] X. Huang and S. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1510–1519, Oct 2017.
- [5] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 694–711. Springer, 2016.
- [6] D. Y. Park and K. H. Lee. Arbitrary style transfer with style-attentional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5880–5888, 2019.
- [7] A. Sanakoyeu, D. Kotovenko, S. Lang, and B. Ommer. A style-aware content loss for real-time hd style transfer. In *European Conference on Computer Vision (ECCV)*, pages 698–714. Springer, 2018.
- [8] N. Van Noord, E. Hendriks, and E. Postma. Toward discovery of the artist's style: Learning to recognize artists by their artworks. *IEEE Signal Processing Magazine*, 32(4):46–54, 2015.
- [9] J. Yang, L. Chen, L. Zhang, X. Sun, D. She, S.-P. Lu, and M.-M. Cheng. Historical context-based style classification of painting images via label distribution learning. In ACM Multimedia, pages 1154–1162, 2018.
- [10] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, Oct 2017.