
Disentangling Style and Content in Anime Illustrations

Sitao Xiang¹

Hao Li^{1,2,3}

¹University of Southern California

²Pinscreen

³USC Institute for Creative Technologies

sitaoxia@usc.edu

hao@hao-li.com

Abstract

Existing methods for AI-generated artworks still struggle with generating high-quality stylized content, where high-level semantics are preserved, or separating fine-grained styles from various artists. We propose a novel Generative Adversarial Disentanglement Network which can fully decompose complex anime illustrations into style and content. Training such model is challenging, since given a style, various content data may exist but not the other way round. In particular, we disentangle two complementary factors of variations, where one of the factors is labelled. Our approach is divided into two stages, one that encodes an input image into a style independent content, and one based on a dual-conditional generator. We demonstrate the ability to generate high-fidelity anime portraits with a fixed content and a large variety of styles from over a thousand artists, and vice versa, using a single end-to-end network and with applications in style transfer. We show this unique capability as well as superior output to the current state-of-the-art.

1 Introduction

Computer generated art [8] has become a topic of focus lately, due to revolutionary advancements in deep learning. Neural style transfer [5] is a groundbreaking approach where high-level styles from artwork can be re-targeted to photographs using deep neural networks. While there has been numerous works and extensions on this topic, there are deficiencies in existing methods. For complex artworks, the methods that rely on matching neural network features and feature statistics, do not sufficiently capture the concept of style at the semantic level. Methods based on image-to-image translation [11] are able to learn domain specific definitions of style, but do not scale well to a large number of styles.

In addressing these challenges, we found that style transfer can be formulated as a particular instance of a general problem, where the dataset has two complementary factors of variation, with one of the factors labelled, and the goal is to train a generative network where the two factors can be fully disentangled and controlled independently. For the style transfer problem, these two factors are style and content.

Based on various adversarial training techniques, we propose a solution to the problem of these two factors and call our method Generative Adversarial Disentangling Network. Our approach consists of two main stages. First, we train a style-independent content encoder, then we introduce a dual-conditional generator based on auxiliary classifier GANs. We demonstrate the disentanglement performance of our approach on a large dataset of anime portraits with over a thousand artist-specific styles, where our decomposition approach outperforms existing methods in terms of level of details and visual quality. Our method can faithfully generate portraits with proper style-specific aspect ratios, shapes, and appearances of facial features, including eyes, mouth, chin, hair, blushes, highlights,

contours, as well as overall color saturation and contrast. In particular, we show how artist-specific styles can be recognized even though the results were synthesized artificially.

2 Background

Neural Style Transfer. Gatys et al. [5] proposed a powerful idea of decomposing an image into content and style using a deep neural network. By matching the features extracted from one input image using a pre-trained network and the Gram matrix of features of another image, one can optimize for an output image that combines the content of the first input and the style of the second one. One area of improvement has been to first identify the best location for style source in the style image for every location in the content image. Luan et al. [18] use masks from either user input or semantic segmentation for guidance. Liao et al. [16] extends this approach by finding dense correspondences between the images using deep neural network features in a coarse-to-fine fashion. In another line of research, different representations for style have been proposed. In Huang et al. [9], the style is represented by affine transformation parameters of instance normalization layers. These methods go slightly beyond transferring texture statistics. While we agree that texture features are an important part of style, these texture statistics do not capture high-level style semantics. Our observation is that, the concept of “style” is inherently domain-dependent, and no set of features defined a priori can adequately handle style transfer problem in all domains at once, let alone mere texture statistics. In our views, a better way of addressing style transfer would be to pose it as an image-to-image translation problem.

Image-to-image Translation. Isola et al. introduced the seminal work on image-to-image translation in [11] where extensive source and target training data need to be supplied. Several extensions to this approach, such as CycleGAN [25] and DualGAN [24], removed the need for supervised training, which significantly increases the applicability of such methods. Of particular interest, Zhu et al. [25] demonstrates the ability to translate between photorealistic images and the styles of Van Gogh, Monet and Ukiyo-e. While the results are impressive, a different network is still required for each pair of domains of interest. One solution to this, is to train an encoder and a generator for each domain such that each domain can be encoded to and generated from a shared code space, as described in Liu et al. [17]. We wish to take one step further and use only one set of networks for many domains: the same encoder encodes each domain into a common code space, and the generator receives domain labels that can generate images for different domains. In a way, rather than considering each style as a separate domain, we can consider the whole set of images as a single large domain, with content and style being two different factors of variation, and our goal is to train an encoder-decoder that can disentangle these two factors. The idea of supplying domain information so that a single generator can generate images in different domains is also used in StarGAN [2]. But in their work, there is no explicit content space, so it does not achieve the goal of disentangling style and content. Our view of style transfer being an instance of disentangled representation is also shared by [10], but their work considers mapping between two domains only.

Disentangled Representation. DC-IGN [15] achieves clean disentanglement of different factor of variation in a set of images. However, the method requires very well-structured data. In particular, the method requires batches of training data with the same content but different style, as well as data with the same style but different content. For style transfer, it is often impossible to find many or even two images that depict the exact same content in different styles. On the other extreme, the method of Chen et al. [1] is unsupervised and can discover disentangled factors of variations from unorganized data. Being unsupervised, there is no way to enforce the separation of a specific set of factors explicitly. The problem we are facing is in between, as we would like to enforce the meaning of the disentangled factors (style and content), but only one of the factors is controlled in the training data, as we can find images presumed to be in the same style that depict different content, but not vice versa. [20] is one example where the setting is the same to ours. Interestingly, related techniques can also be found in the field of audio processing. The problem of voice conversion, where an audio speech is mapped to the voice of a different speaker, has a similar structure to our problem. In particular, our approach is similar to [3].

3 Method

The training data must be organized by style. However, fine-grained labels of styles are difficult to obtain. Hence, we use the identity of artists, as a proxy for style. While an artist might have several styles and it might evolve over time, using an artist’s identity as a proxy for style is a good approximation and an efficient choice, since the artist’s label are readily available. As in [3], our method is divided into two stages.

Stage 1: Style Independent Content Encoding. In this stage, the goal is to train an encoder that encode as much information as possible about the content of an image, but no information about its style. We use per pixel L2 distance (not its square) for the reconstruction loss: for two (3-channel) images $X, Y \in \mathbb{R}^{h \times w \times 3}$, by $\|X - Y\|$ without any subscript we mean this distance:

$$\|X - Y\| = \frac{1}{hw} \sum_{i=1}^h \sum_{j=1}^w \|X_{ij} - Y_{ij}\|_2$$

We start from a solution that did not work well. Consider a simple encoder-decoder network with encoder $E(\cdot)$ and decoder $G(\cdot)$ whose sole purpose is to minimize the reconstruction loss:

$$\mathcal{L}_{\text{rec}} = \mathbb{E}_{x \sim p(x)} [\|x - G(E(x))\|] \quad \min_{E, G} \mathcal{L}_{\text{rec}}$$

where $p(x)$ is the distribution of training samples. To prevent The encoder from encoding style information, we add an adversarial classifier $C(\cdot)$ that tries to classify the encoder’s output by artist, while the encoder tries to maximize the classifier’s loss:

$$\begin{aligned} \mathcal{L}_C = \mathbb{E}_{x, a \sim p(x, a)} [\text{NLL}(C(E(x)), a)] \quad \min_C \mathcal{L}_C \\ \min_{E, G} \mathcal{L}_{\text{rec}} - \lambda \mathcal{L}_C \end{aligned} \quad (1)$$

where a is the integer ground truth label representing the author of image x and $p(x, a)$ is the joint distribution of images and their authors. λ is a weight factor. To resolve the conflicting goal that the generator needs style information to reconstruct the input but the encoder must not encode style information to avoid being successfully classified, we can learn a vector for each artist, representing their style. This “style code” is then provided to $G(\cdot, \cdot)$ which now takes two inputs. We introduce the style function $S(\cdot)$ which maps an artist to their style vector. Now the objective is

$$\begin{aligned} \mathcal{L}_{\text{rec}} = \mathbb{E}_{x, a \sim p(x, a)} [\|x - G(E(x), S(a))\|] \\ \min_{E, G, S} \mathcal{L}_{\text{rec}} - \lambda \mathcal{L}_C \end{aligned}$$

Note that $S(\cdot)$ is not another encoding network. It does not see the input image, but only its artist label. This is essentially identical to the first stage of [3]. In our experiments, we found that this method does not adequately prevent the encoder from encoding style information. We discuss this issue in appendix B.

We propose the following changes: instead of the code $E(x)$, $C(\cdot)$ tries to classify the generator’s output $G(E(x), S(a'))$, which is the combination of the content of x and the style of a *different* artist a' . In addition, akin to Variational Autoencoders [14], the output of $E(\cdot)$ and $S(\cdot)$ are parameters of multivariate normal distributions and we introduce KL-divergence loss to constrain these distributions. To avoid the equations becoming too cumbersome, we overload the notation a bit so that when $E(\cdot)$ and $S(\cdot)$ are given to another network as input we implicitly sample a code from the distribution. The optimization objectives becomes

$$\begin{aligned} \mathcal{L}_C = \mathbb{E}_{\substack{x, a \sim p(x, a) \\ a' \sim p(a)}} [\text{NLL}(C(G(E(x), S(a'))), a)] \quad \min_C \mathcal{L}_C \\ \mathcal{L}_{\text{rec}} = \mathbb{E}_{x, a \sim p(x, a)} [\|x - G(E(x), S(a))\|] \\ \mathcal{L}_{E\text{-KL}} = \mathbb{E}_{x \sim p(x)} [D_{\text{KL}}(E(x) || \mathcal{N}(0, I))] \quad \mathcal{L}_{S\text{-KL}} = \mathbb{E}_{a \sim p(a)} [D_{\text{KL}}(S(a) || \mathcal{N}(0, I))] \\ \min_{E, G, S} \mathcal{L}_{\text{rec}} - \lambda_C \mathcal{L}_C + \lambda_{E\text{-KL}} \mathcal{L}_{E\text{-KL}} + \lambda_{S\text{-KL}} \mathcal{L}_{S\text{-KL}} \end{aligned}$$

Stage 2: Dual-Conditional Generator. It is well known that autoencoders typically produce blurry outputs and fail to capture texture information which is an important part of style. To condition on styles, we base our approach on auxiliary classifier GANs [22]. First, as usual, a discriminator $D(\cdot)$ tries to distinguish between real and generated samples, but instead of binary cross entropy, we follow [19] and use least squares loss. The discriminator’s loss is

$$\mathcal{L}_{D\text{-real}} = \mathbb{E}_{x \sim p(x)} [(D(x) - 1)^2] \quad \mathcal{L}_{D\text{-fake}} = \mathbb{E}_{\substack{x \sim p(x) \\ a' \sim p(a)}} [(D(G(E(x), S(a')))) + 1)^2]$$

while the generator’s loss against the discriminator is

$$\mathcal{L}_{D\text{-adv}} = \mathbb{E}_{\substack{x \sim p(x) \\ a' \sim p(a)}} [D(G(E(x), S(a')))]^2$$

Here the encoder $E(\cdot)$, generator $G(\cdot, \cdot)$ and style function $S(\cdot)$ are inherited from stage 1. Note that in all the equations a is sampled jointly with x but a' is independent. Then, similar to [22], a classifier is trained to classify training images to their correct authors:

$$\mathcal{L}_{C_2\text{-real}} = \mathbb{E}_{x, a \sim p(x, a)} [\text{NLL}(C_2(x), a)]$$

Unlike the generator and encoder, $C_2(\cdot)$ is a different classifier than the one in stage 1, thus we add subscript 2 to disambiguate, and shall refer to the stage 1 classifier as $C_1(\cdot)$. The generator try to generate samples that would be classified as the artist that it is conditioned on, by adding this to its loss function:

$$\mathcal{L}_{C_2\text{-adv}} = \mathbb{E}_{\substack{x \sim p(x) \\ a' \sim p(a)}} [\text{NLL}(C_2(G(E(x), S(a'))), a')] \quad (2)$$

But we differ from previous works on conditional GANs in the treatment of generated samples by the classifier. In [22] the classifier is cooperative in the sense that it would also try to minimize equation 2. In other works like [2, 3, 20] there is no special treatment for generated images, and the classifier does not optimize any loss function on them. In [23] the classifier is trained to be uncertain about the class of generated samples, by maximizing the entropy of $C_2(G(E(x), S(a')))$.

We train the classifier it to explicitly classify generated images conditioned on the style of a as “not a ”. For this, we define what we call the “negative log-likelihood”:

$$\text{NLU}(\mathbf{y}, i) = -\log(1 - y_i)$$

and take this as the classifier’s loss on generated samples:

$$\mathcal{L}_{C_2\text{-fake}} = \mathbb{E}_{\substack{x \sim p(x) \\ a' \sim p(a)}} [\text{NLU}(C_2(G(E(x), S(a'))), a')]$$

We discuss the effect of an adversarial classifier in appendix B. While the discriminator and the classifier are commonly implemented as a single network with two outputs, we use separate networks for $C_2(\cdot)$ and $D(\cdot)$. To enforce the condition on content, we simply take E and require that the generated samples be encoded back to its content input:

$$\mathcal{L}_{\text{cont}} = \mathbb{E}_{\substack{x \sim p(x) \\ a' \sim p(a)}} [\|E(G(E(x), S(a'))) - E(x)\|_2^2]$$

And our training objective for this stage is:

$$\begin{aligned} \min_D \mathcal{L}_{D\text{-real}} + \mathcal{L}_{D\text{-fake}} \quad \min_{C_2} \mathcal{L}_{C_2\text{-real}} + \mathcal{L}_{C_2\text{-fake}} \\ \min_{G, S} \lambda_D \mathcal{L}_{D\text{-adv}} + \lambda_{C_2} \mathcal{L}_{C_2\text{-adv}} + \lambda_{\text{cont}} \mathcal{L}_{\text{cont}} + \lambda_{S\text{-KL}} \mathcal{L}_{S\text{-KL}} \end{aligned}$$

Note that $E(\cdot)$ is fixed in stage 2. Since at the equilibrium $C_2(\cdot)$ is expected to classify real samples effectively, it would be good to pre-train $C_2(\cdot)$ on real samples only prior to stage 2. The training procedures are summarized in figure 1. Any computation marked in red are for discriminators and classifiers only, those marked in blue are for the generator only, and those marked in black are common to both. KL-divergence omitted for clarity.

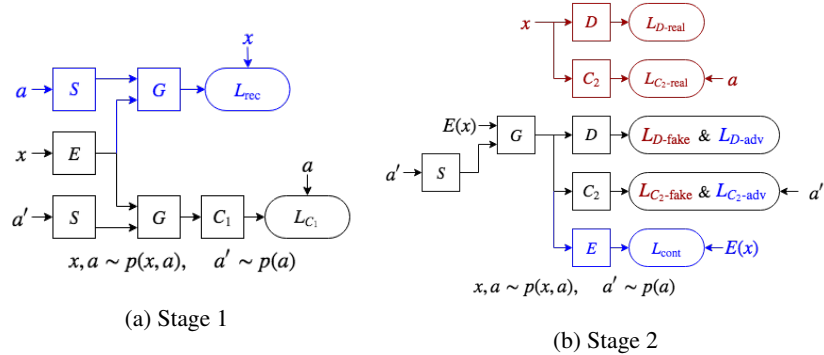


Figure 1: Training procedure

Table 1: Common part of network architecture

Layer	-	SC	SC	SC	SC	SC	SC	F
Channel	3	32	64	128	256	512	1024	2048
Size	256	128	64	32	16	8	4	-

4 Implementation

Dataset. We obtain our training data from Danbooru¹. We took all images with exactly one artist tag, and processed them with an open source tool AnimeFace 2009 [21] to detect faces. Each detected face is rotated into an upright position, then cropped out and scaled to 256×256 pixels. Every artist who has at least 50 image were used for training. Our final training set includes 106,814 images from 1139 artists.

Network Architecture. All networks are built from residue blocks [7]. While the residue branch is commonly composed of two convolution layers, we only use one convolution per block and increase the number of blocks instead. ReLU activations are applied on the residue branch before it is added to the shortcut branch. For simplicity, all of our networks had an almost identical structure, with only minor differences in the input/output layers. The common part is given in table 1. The network is composed of the sequence of blocks in the first row (C = stride 1 convolution, S = stride 2 convolution, F = fully connected) with the number of output channels in the second row and spatial size of output feature map in the third row. For the generator, the sequence runs from right to left, while for other networks the sequence runs from left to right. On top of this common part, fully connected input/output layers are added for each network, with appropriate number of input/output features: for classifiers $C_1(\cdot)$ and $C_2(\cdot)$, the number of artists is 1,139; for the discriminator $D(\cdot)$, we use 1; for the encoder $E(\cdot)$, we use 2 parallel layers with 256 features for mean and standard deviation of output distribution; for the generator $G(\cdot, \cdot)$, the sum of the content and style dimensions is 512.

Training. We weight the loss terms with the hyperparameters in table 2, which also includes training parameters. $S(\cdot)$ is treated differently, since it is not a network but just matrices storing style codes. We use the same learning rate for every network in a single stage. In stage 2, we use RMSprop for our weight updating algorithm since using the momentum would sometimes cause instabilities in GANs. In every other stage we use the Adam optimizer. Timing is measured as number of iterations.

5 Results

Disentangled Representation of Style and Content. To test our main goal of disentangling style and content in the generator’s input code, we show that the style code and content code do indeed only control the respective aspect of the generated image, by fixing one and changing the other. In figure 2, in each group of two rows, the leftmost two images are examples of illustrations drawn by

¹danbooru.donmai.us

Table 2: Weighting and training hyperparameters

Weight	Value						
		Stage	Learning rate		Algorithm	Batch	Time
				S	Others		
λ_{C_1}	0.2	1	0.005	5×10^{-5}	Adam	8	400k
λ_{E-KL}	10^{-4}	C_2 pre-train		-	Adam	16	200k
λ_{S-KL}	2×10^{-5}	2	0.01	2×10^{-5}	RMSprop	8	400k
λ_D	1						
λ_{C_2}	1						
λ_{cont}	0.05						



Figure 2: Images generated from fixed style and different contents, for two artists. Leftmost column taken from training set, courtesy of respective artists. Top group: Sayori. Bottom group: Swordsouls.



Figure 3: Images generated from the same content and different styles. Contain both style of artists from the training set and style codes randomly samples from the style distribution

an artist in the dataset, and to the right are 12 samples generated from the style of the artist. Since this can be expected from a conventional class-conditional GAN, it is not the main contribution of our method. In particular, the strength of our method, is the ability to generate the same content in different styles where facial shapes, appearances, and aspect ratios are faithfully captured. We refer to Appendix A for additional results and discussions. Figure 3 shows 35 images generated with a fixed content and different style codes.

Style Transfer. As a straightforward application, we show some style transfer results in figure 4, with comparisons to existing methods, the original neural style transfer [5] and StarGAN [2]. We can see that neural style transfer seems to mostly apply the color of the style image to the content image, which, by our definition, not only fails to capture the style of the style image but also alters the content of the content image. StarGAN managed to transfer the overall use of color of the target artist and some other prominent features like size of the eyes, but otherwise fails to capture intricate style elements. Our method transfers the style of the target artist much more faithfully.

Evaluation. Other than visual results, there is no well-established quantitative measure for the quality of style transfer methods and accessing experts for anime styles is challenging for a proper user study. The evaluation in [3] is audio specific. In [2], the classification accuracy on the generated samples is given as a quality measure. However, this seemingly reasonable measure, as we shall argue in appendix B.2, is not adequate. Nevertheless, we do report that the samples generated by our generator can be classified by style by our classifier with 86.65% top-1 accuracy. Detailed testing procedure, results and analysis can be found in appendix B.2. A comprehensive ablation study is given in Appendix B, where major differences from previous approaches are evaluated. We have also approached several anime enthusiasts and they were able to distinguish the synthesized styles based on the artworks of known artists.

Limitations. Due to the relative scarcity of large labelled collection of artworks depicting a consist set of subjects, we have thus far not collected and tested our method on examples beyond portraits and artistic styles beyond anime illustrations. We also noticed some inconsistencies in small features, such as eye colors, presumably due to small features being disfavored by the per-pixel reconstruction loss in stage 1, as well as our choice of architecture with fixed-size code rather than fully convolutional. Additional discussions can be found in Appendix A.

6 Conclusion

We introduced a Generative Adversarial Disentangling Network which enables true semantic-level artwork synthesis using a single generator. Our evaluations and ablation study indicate that style and content can be disentangled effectively through our a two-stage framework, where first a style independent content encoder is trained and then, a content and style-conditional GANs is used for synthesis. While we believe that our approach can be extended to a wider range of artistic styles, we have validated our technique on various styles within the context of anime illustrations. In particular, this techniques is applicable, as long as we disentangle two factors of variation in a dataset and only one of the factors is labelled and controlled. Compared to existing methods for style transfer, we show significant improvements in terms of modeling high-level artistic semantics and visual quality. We found that anime experts were able to identify artist-specific styles correctly in our artificially generated illustrations. In the future, we hope to extend our method to styles beyond anime artworks, and we are also interested in learning to model entire character bodies, or even entire scenes.

Acknowledgments

We thank all the illustrators who created these magnificent artworks that made this project possible. In particular, Sitao Xiang wishes to dedicate this work to Sayori - his favorite illustrator - and her catgirls.

Hao Li is affiliated with the University of Southern California, the USC Institute for Creative Technologies, and Pinscreen. This research was conducted at USC and was funded by in part by the ONR YIP grant N00014-17-S-FO14, the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, the Andrew and



Figure 4: Style transfer results. In the top two rows, in each column are two samples from the training dataset by the same artist. In each subsequent group of three rows, the leftmost image is from the training dataset. The images to the right are style transfer results generated by three different methods, from the content of the left image in the group and from the style of the top artist in the column. In each group, first row is our method, second row is StarGAN and third row is neural style. For neural style, the style image is the topmost image in the column. Training samples courtesy of respective artists. Style samples, from left: Sayori, Ideolo, Peko, Yabuki Kentarou, Coffee-kizoku, Mishima Kurone, Ragho no Erika. Content samples, from top: Kantoku, Koi, Horiguchi Yukiko.

Erna Viterbi Early Career Chair, the U.S. Army Research Laboratory (ARL) under contract number W911NF-14-D-0005, Adobe, and Sony. This project was not funded by Pinscreen, nor has it been conducted at Pinscreen or by anyone else affiliated with Pinscreen. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

- [1] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016.
- [2] Yunje Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, 2018.
- [3] Ju-chieh Chou, Cheng-chieh Yeh, Hung-yi Lee, and Lin-shan Lee. Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations. *arXiv preprint arXiv:1804.02812*, 2018.
- [4] Leon Gatys, Alexander S Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 262–270, 2015.
- [5] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.
- [6] Leon A Gatys, Alexander S Ecker, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. Controlling perceptual factors in neural style transfer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [8] Aaron Hertzmann. Can computers create art? *CoRR*, abs/1801.04486, 2018. URL <http://arxiv.org/abs/1801.04486>.
- [9] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 1510–1519. IEEE, 2017.
- [10] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018.
- [11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976. IEEE, 2017.
- [12] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*, 2018.
- [13] Hadi Kazemi, Seyed Mehdi Iranmanesh, and Nasser Nasrabadi. Style and content disentanglement in generative adversarial networks. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 848–856. IEEE, 2019.
- [14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [15] Tejas D Kulkarni, William F Whitney, Pushmeet Kohli, and Josh Tenenbaum. Deep convolutional inverse graphics network. In *Advances in neural information processing systems*, pages 2539–2547, 2015.

- [16] Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. Visual attribute transfer through deep image analogy. *ACM Transactions on Graphics (TOG)*, 36(4):120, 2017.
- [17] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems*, pages 700–708, 2017.
- [18] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6997–7005. IEEE, 2017.
- [19] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2813–2821. IEEE, 2017.
- [20] Michael F Mathieu, Junbo Jake Zhao, Junbo Zhao, Aditya Ramesh, Pablo Sprechmann, and Yann LeCun. Disentangling factors of variation in deep representation using adversarial training. In *Advances in Neural Information Processing Systems*, pages 5040–5048, 2016.
- [21] nagadomi. Animeface 2009. <https://github.com/nagadomi/animeface-2009>, 2017. Accessed: 2018-1-9.
- [22] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International Conference on Machine Learning*, pages 2642–2651, 2017.
- [23] Jost Tobias Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. *arXiv preprint arXiv:1511.06390*, 2015.
- [24] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2868–2876. IEEE, 2017.
- [25] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2242–2251. IEEE, 2017.

A Additional Discussions

A.1 Thoughts on Style and Style Transfer

While works on this topic has been numerous, we feel that one fundamental question is not often carefully addressed: what is style, in a deep learning setting?

As stated in [5], which is based on an earlier work on neural texture synthesis [4], the justification for using Gram matrices of neural network features as a representation of style is that it captures statistical texture information. So, in essence, “style” defined as such is a term for “texture statistics”, and the style transfer is limited to texture statistics transfer. Admittedly, it does it in smart ways, as in a sense the content features are implicitly used for selecting which part of the style image to copy the texture statistics from.

As discussed in section 2 above, we feel that there is more about style than just feature statistics. Consider for example the case of caricatures. The most important aspects of the style would be what facial features of the subjects are exaggerated and how they are exaggerated. Since these deformations could span a long spatial distance, they cannot be captured by local texture statistics alone.

Another problem is domain dependency. Consider the problem of transferring or preserving color in style transfer. If we have a landscape photograph taken during the day and want to change it to night by transferring the style from another photo taken during the night, or if we want to change the season from spring to autumn, then color would be part of the style we want to transfer. But if we have a still photograph and want to make it an oil painting, then color is part of the content, we may want only the quality of the strokes of the artwork but keep the colors of our original photo.

People are aware of this problem and in [6], two methods, luminance-only style transfer and color histogram matching, are developed to optionally keep the color of the content image. However, color is only one aspect of the image for which counting it as style vs. content could be an ambiguity. For more complicated aspects, the option to keep or to transfer may not be easily available.

We make two observations here. First, style must be more than just feature statistics. Second, the concept of “style” is inherently domain-dependent. In our opinion, “style” means different ways of presenting the same subject. In each different domain, the set of possible subjects is different and so is the set of possible ways to present them.

So, we think that any successful style transfer method must be adaptive to the intended domain and the training procedure must actively use labelled style information. Simple feature based methods will never work in the general setting. This includes previous approaches which explicitly claimed to disentangle style and content, such as in [13] which adopts the method in the original neural style transfer for style and content losses, and also some highly accomplished methods like [16].

As a side note, for these reasons we feel that some previous methods made questionable claims about style. In particular, works like [10] and StyleGAN [12] made reference to style while only being experimented on collections of real photographs. By our definition, in such dataset, without a careful definition and justification there is only one possibly style, that is, photorealistic, so the distinction between style and content does not make sense, and calling a certain subset of factors “content” and others “style” could be an arbitrary decision.

This is also why we elect to not test our method on more established GAN datasets like CelebA or LSUN, which are mostly collections of real photos.

A.2 More on Results

The differences in style can be subtle, and readers may not be familiar enough with anime illustrations to be able to recognize them. We hint at several aspects to which the reader may want to pay attention: overall saturation and contrast, prominence of border lines, overall method of shading (flat vs. 3-dimensional), size and aspect ratio of the eyes, shape of the irides (round vs. angular), shininess of the irides, curvature of the upper eyelids, height of lateral angle of the eyes, aspect ratio of the face, shape of the chin (round vs. angular), amount of blush, granularity of hair strands, prominence of hair shadow on forehead, treatment of hair highlight (intensity, no specularities / dots along a line / thin line / narrow band / wide band, clear smooth boundary / jagged boundary / fuzzy boundary, etc.).

If we specifically examine these style elements in the style transfer results in figure 4, it should be evident that our method has done a much better job. That being said, our network is not especially good at preserving fine content details.

This can be said to be a result from our choice of architecture: we use an encoder-decoder with a fixed-length content code while StarGAN uses a fully convolutional generator with only 2 down-sampling layers, which is able to preserve much richer information about the content image.

Part of the reason is that we consider it a desirable feature to be able to sample from the content distribution. In convolutional feature maps, features at different locations are typically correlated in complicated ways which makes sampling from the content distribution impractical. Unsurprisingly previous works adopting fully convolutional networks did not demonstrate random sampling of content, which we did.

Another reason is that we feel that the ability of fully convolutional networks to preserve content has a downside: it seems to have a tendency to preserve image structures down to the pixel level. As can be observed in figure 4, StarGAN almost exactly preserves location of border lines between different color regions. As we have mentioned, part of the artists’ style is different shape of facial features, and a fully convolutional network struggle to capture this. We suspect that for a fully convolutional architecture to work properly, some form of non-rigid spatial transformation might be necessary. This can be one of our future directions.

A.3 Inconsistency of Small Features

As can be seen from many of the visualizations, the color of the eyes, and facial expression - which largely depends on the shape of the mouth - are sometimes not well preserved when changing the style codes. We thought that this could be the result of the choice of the reconstruction loss function in stage 1. The loss is averaged across all image pixels, which means the importance of a feature to the network is proportional to its area. With the code length being limited, the network prioritizes encoding large features, which can be seen from the fact that large color regions in the background are more consistent than eyes and mouth.

For this particular problem, we can utilize additional tags: almost all images gathered from Danbooru are tagged according to eye color and facial expressions. We could train the encoder to classify for these tags while the generator could be conditioned on these tags. For general problems however, we may need loss functions that are more aligned with humans’ perception of visual importance.

Another observation is that sometimes in a generated image the two eyes are of different color. We point out that the dataset do contain such images, and while still rare, the prevalence of heterochromia in anime characters is much higher than in real life. We think that this could be enough to force the encoder to use two different set of dimensions to encode the color of each eye, thus causing random samples to have different colored eyes.

B Ablation Study

We provide additional details on our method in Section 3, along with ablation studies to demonstrate their effectiveness.

B.1 Using Reconstruction Result as Stage 1 Classifier Input

By design, assuming sufficiently powerful networks, the stage 1 method in [3] should reconstruct x perfectly with $G(E(x), S(a))$ while at the same time not include any style information in $E(x)$. But in reality, this method worked poorly in our experiments. We discuss our conjecture here.

We think the problem is that the distribution of $E(x)$ is unconstrained and the encoder-decoder can exploit this to trick the classifier while still encode style information in $E(x)$. Assume that the last layer weight of $E(\cdot)$ is $W_1 \in \mathbb{R}^{m \times d}$ and the first layer weight of $G(\cdot, \cdot)$ (connected to its content input) is $W_2 \in \mathbb{R}^{d \times n}$ where d is the length of the content code and n and m are the number of features in relevant layers. Then for any invertible matrix $H \in \mathbb{R}^{d \times d}$, replacing W_1 with $W_1 H$ and W_2 with $H^{-1} W_2$ would give us an autoencoder that computes the exact same function as before but with a linearly transformed distribution in the code space. Note that such a “coordination” is not




Figure 5: Comparison of stage 1 image reconstruction with correct style and zero style, using different methods. Column 1: images from the dataset. Column 2: VAE reconstruction. Column 3: MLP classifier, correct style. Column 4: MLP classifier, zero style. Column 5: Our classifier, correct style. Column 6: Our classifier, zero style. Training images courtesy of respective artists. From top: Azumi Kazuki, Namori, Iizuki Tasuku, Tomose Shunsaku.

limited to be between the innermost layers, and the parameters of $E(\cdot)$ and $G(\cdot, \cdot)$ may be altered to keep the reconstruction almost unchanged but at the same time transform the code distribution in complicated ways beyond simple linear transformation.

As a result, the encoder-decoder can transform the code distribution constantly during training. The classifier would see a different input distribution in each iteration and thus fail to learn to infer the artist from style information that could possibly be included in the output of $E(\cdot)$, and thus would be ineffective at preventing $E(\cdot)$ from encoding style information.

Note that constraining the code distribution of the whole training set does not help. For example, if we constraint the code distribution to be standard normal as in VAE, we can take H in the previous example to be an orthonormal matrix. The transformed distribution on the whole dataset would still be standard normal but the distribution of artworks by each different artist would change.

We noticed that this problem would be alleviated if instead of the code $E(x)$, $C_1(\cdot)$ tries to classify the reconstruction result $G(E(x), S(a))$ since the output distribution of the decoder is “anchored” to the distribution of training images and cannot be freely transformed. But the output of $G(\cdot, \cdot)$, being an reconstruction of the input image, will inevitably contain some style information, so giving $G(E(x), S(a))$ as input to $C_1(\cdot)$ will again cause the problem of conflicting goals between reconstruction and not encoding style information. The important observation is that, the input given to $C_1(\cdot)$ is allowed to contain style information as long as it is not the style of the real author of x ! Thus we arrived at our proposed method where the classifier sees $G(E(x), S(a'))$, the image generated by the encoder’s output on x and the style of an artist a' different than the author of x .

Strictly speaking, by classifying the output of the generator we can only ensure that $E(\cdot)$ and $G(\cdot, \cdot)$ would jointly not retain the style of x in $G(E(x), S(a'))$. $E(\cdot)$ may still encode style information which is subsequently ignored by $G(\cdot, \cdot)$. To solve this, we constrain the output distribution of $E(\cdot)$ with a KL-divergence loss as in VAE. With such a constraint, $E(\cdot)$ would be discouraged from encoding information not used by $G(\cdot, \cdot)$ since that would not improve the reconstruction but only increase the KL-divergence penalty.

To compare the effect of classifiers seeing different inputs, we train an alternative version of stage 1 using the method in [3], where instead of a convolutional network that gets the generator’s output, we

use a multi-layer perceptron that gets the encoder’s output as the classifier. We use an MLP with 10 hidden layers, each having 2048 features. We then compare $G(E(x), \mathbf{0})$ with $G(E(x), S(a))$, that is, image generated with the content code of an image x plus an all-zero style code versus the correct style code. A successful stage 1 encoder-decoder should reconstruct the input image faithfully with the correct style code but generate style-neutral images with the zero style code. We also give the reconstruction of an conventional VAE for reference. The result is shown in figure 5.

As autoencoders are in general bad at capturing details, the readers may want to pay attention to more salient features like size of the eyes, shape of facial contour and amount of blush. We can see that these traits are clearly preserved in the style-neutral reconstruction when an MLP is used for classifying $E(x)$. Our method, while not able to completely deprive the output of style information from the input in the style-neutral reconstruction and also altered the content by a bit, performed much better, and did so without compromising the reconstruction with the correct style.

We point out that this is not due to the MLP not having enough capacity to classify $E(x)$: when the encoder is trained to only optimize for reconstruction and ignore the classifier by setting $\lambda_{C_1} = 0$, the classifier is able to classify $E(x)$ with higher than 99% top-1 accuracy.

When setting $\lambda_{C_1} = 0.2$ as normal, the top-1 accuracy drops to a mere 1.52%. While this is not a definitive proof of our conjecture, the behaviour that the encoder is able to trick the classifier while still sneaking in style information is consistent with the scenario we described.

B.2 Adversarial Stage 2 Classifier

One important difference in our style-conditional GAN training compared to previous class-conditional GAN approaches is that, in our case, not only is the discriminator adversarial but the classifier as well. The classifier explicitly tries to classify samples generated using an artist a as “not by a ”.

The rationale is that, by just correctly classifying the real samples, the classifier may not be able to learn all the aspects about an artist’s style. Imagine the hypothetical scenario where each artist draws the hair as well as the eyes differently. During training, the classifier arrives at a state where it can classify the real samples perfectly by just considering the hair and disregarding the eyes. There will then be no incentive for the classifier to learn about the style of the eyes since that will not improve its performance.

If the classifier is adversarial, this will not happen: as long as the generator has not learned to generate the eyes in the correct style, the classifier can improve by learning to distinguish the generated samples from real ones by the style of their eyes.

In general, a non-adversarial classifier only needs to learn as much as necessary to tell different artists apart, while an adversarial classifier must understand an artist’s style comprehensively.

To study the effect of this proposed change, we repeat stage 2 but without adding $\mathcal{L}_{C_2\text{-fake}}$ to the classifier’s loss.

In figure 6 we generate samples using this generator trained without an adversarial classifier, from the exact same content code and artists as in figure 2. While we leave the judgment to the readers, we do feel that making the classifier adversarial does improve the likeness.

The choice of “negative log-likelihood” as the classifier’s loss on generated samples might seem a bit unconventional:

$$\begin{aligned} \text{NLU}(\mathbf{y}, i) &= -\log(1 - y_i) \\ \mathcal{L}_{C_2\text{-fake}} &= \mathbb{E}_{\substack{x \sim p(x) \\ a' \sim p(a)}} [\text{NLU}(C_2(G(E(x), S(a'))), a')] \end{aligned}$$

To many, maximizing the negative log-likelihood would be a much more natural choice. The major concern here is that the negative log-likelihood is not upper-bounded. Suppose that $C_2(\cdot)$ is trained using such an objective. There would then be little to be gained by trying to classify real samples correctly because $\mathcal{L}_{C_2\text{-real}}$ is lower-bounded. To decrease its loss, it would be much more effective to increase the negative log-likelihood on generated samples. The result is that the $C_2(\cdot)$ simply ignores the loss on real samples while the negative log-likelihood on generated samples quickly explodes. If



Figure 6: Images generated from fixed style and different contents, when stage 2 classifier is not adversarial.

a classifier cannot classify the real samples correctly, one would not expect the generator to learn the correct styles by training it against that classifier.

In contrast, in stage 1 when we trained the encoder $E(\cdot)$ against the classifier $C_1(\cdot)$, we maximize the negative log-likelihood. The same problem does not occur here because it is $E(\cdot)$ and $G(\cdot, \cdot)$, not $C_1(\cdot)$, who tries to maximize it. Considering that the pixel values of the generator’s output has a finite range, for a fixed classifier, the log-likelihood is bounded.

Indeed, in stage 1 the negative log-unlikelihood would not be a favourable choice: in the equilibrium, the classifier is expected to perform poorly (no better than random guess), at which point the negative log-unlikelihood would have tiny gradients while the negative log-likelihood would have large gradients.

A Note on Using Classification Accuracy as Quality Measure. From the discussion above, we conclude that the accuracy of classification by style on the generated samples with an independently trained classifier cannot properly assess the level of success of style transfer: since such a classifier, trained solely on real samples, cannot capture complete style information, a high accuracy achieved by generated samples does not indicate that the style transfer is successful. It is likely, though, that the generator is of bad quality if the accuracy is very low.

However, we feel that if the generator could score a high accuracy even against an adversarial classifier, it must be considered a success. Unfortunately, our method has yet to achieve this goal.

In table 3 we report the classification accuracy by both the adversarial and the non-adversarial classifier on samples generated both by the generator trained against the adversarial classifier and by the generator trained against the non-adversarial classifier. Given that for the majority of artists the number of artworks available is quite limited, to use the data at maximum efficiency we used all data for training and did not have a separate test set. As a substitute, we generate samples from content codes drawn randomly from the learned content distribution which is modeled by a multivariate normal distribution over the whole training set. Notice that during stage 2 training we only use $E(x)$ where x is from the training set and $E(\cdot)$ is fixed, so the network only ever sees a fixed finite set of different content codes. Thus, random drawn content codes are indeed unseen which serves the same purpose as a separate test set.

In each test, the same number of samples as in the training set (106,814) are generated such that the content codes are independent and random and each artist appear exactly as many times as in the training set.

Table 3: Top-1 classification accuracy of two classifiers on samples generated from two generators

	Adversarial C_2	Non-adversarial C_2
G trained with adversarial C_2	14.37%	86.65%
G trained with non-adversarial C_2	1.85%	88.59%



Figure 7: Images generated from fixed style and different contents, when explicit condition on content is removed.

As can be seen, there is a huge discrepancy between the classification accuracy scored by the different classifiers. If trained to do so, the adversarial classifier can easily tell that the generated samples are not of the correct style, so extreme caution must be exercised if any conclusions is to be drawn from the classification accuracy by a non-adversarial classifier.

B.3 Explicit Conditioning on Content in Stage 2 Generator

Our treatment of content in stage 2 is also quite different, in that we use a content loss to explicitly condition the image generation on content. Compare this to [3] where there is no explicit condition on content, rather, the stage 2 generator’s output is added to the stage 1 decoder so that the stage 1 condition on content, which is guaranteed by minimizing the reconstruction loss, is inherited by stage 2.

While we did not experiment with their approach, we show here that in our framework the explicit condition on content by content loss is necessary. We repeat stage 2 training with $\lambda_{\text{cont}} = 0$ so that $E(G(E(x), S(a')))$ is not constrained to be close to $E(x)$. We then examine if style and content can still be controlled independently. Images generated by fixing the style and varying the content is shown in figure 7, while images generated by fixing the content and varying the style is shown in figure 8.

We can clearly see that the content code has lost most of its ability to actually control the content of the generated image, while the style code would now control both style and content. Curiously, the content code still seems to control the head pose of the generated character, as this is the most noticeable variation in figure 7 while in figure 8 the characters had essentially the same head pose.

This can be considered a form of “partial mode collapse”: while the network is still able to generate samples with rich variation, some of the input variables have little effect on the generated image, and part of the input information is lost. This happens despite the fact that the stage 2 generator is initialized with the weights of the final stage 1 generator, in which the content code is able to control the content. So, this is not the case where the number of input variables is more than the dimension of the distribution of the dataset so that some input variables never got used. Rather, the ability to control content is initially present but subsequently lost. So, in our approach, an explicit condition on the content by a content loss is necessary.



Figure 8: Images generated from fixed content and different styles, when explicit condition on content is removed.