

# Semi-Supervised Image-to-Image Translation

Manan Oza

*Department of Computer Engineering  
D. J. Sanghvi College of Engineering  
Mumbai, India  
manan.oza0001@gmail.com*

Himanshu Vaghela

*Department of Computer Engineering  
D. J. Sanghvi College of Engineering  
Mumbai, India  
himanshuvaghela1988@gmail.com*

Prof. Sudhir Bagul

*Department of Computer Engineering  
D. J. Sanghvi College of Engineering  
Mumbai, India  
Sudhir.Bagul@djsce.ac.in*

**Abstract**—Image-to-image translation is a long-established and a difficult problem in computer vision. In this paper we propose an adversarial based model for image-to-image translation. The regular deep neural-network based methods perform the task of image-to-image translation by comparing gram matrices and using image segmentation which requires human intervention. Our generative adversarial network based model works on a conditional probability approach. This approach makes the image translation independent of any local, global and content or style features. In our approach we use a bidirectional reconstruction model appended with the affine transform factor that helps in conserving the content and photorealism as compared to other models. The advantage of using such an approach is that the image-to-image translation is semi-supervised, independent of image segmentation and inherits the properties of generative adversarial networks tending to produce realistic. This method has proven to produce better results than Multimodal Unsupervised Image-to-image translation.

**Index Terms**—GANs, image-to-image translation, style transfer

## I. INTRODUCTION

Image-to-image transfer has established itself as an important domain in computer vision since the first paper published by Gatys et al. [1]. Also known as Neural Style Transfer, it has had many variations over the years, image colorization [14], style transfer [1], image-to-image transfer [3] and so on. For which generally deep neural networks have been used with architectural variances. For instance, we can make a day time image (also known as the content image) of a city look like a night time image by selecting the appropriate style (reference) image. Likewise we can have diverse types of features transferred from one image to another which include time, color, seasonal translations as well.

Image-to-image translation is the process of translating one image onto another while preserving the content and photorealism of the original content image. Deep-learning techniques have proved excellent in faithful and photorealistic style translation [1]–[3], [7]. Our approach is built upon the idea of generative adversarial networks introduced by Goodfellow et al. [8]. The underlying concept of such a neural network architecture is that a GAN consists of a generator and a discriminator. The discriminator is trained to identify real images while the generator tries to fool the discriminator by creating counterfeit images from noise and passes them on to the discriminator. Which then returns a verdict on how close the counterfeit images are to a real one. Based on this feedback

the generator improves itself and creates another image and the cycle repeats.

Here in our paper we make use of an improvised GAN architecture appended with an Affine Loss factor calculated from a Matting Laplacian matrix [6] in the final loss function. This additional factor helps in maintaining spatial integrity and preserve photorealism in the content image. Since generative adversarial networks create images from noise they are prone to distortions and noisy images but provide with the biggest advantage, they do not form the basis of simple color and style mapping. They recreate the content image with the style variations.

## II. RELATED WORK

Image-to-image style transfer has reached state-of-the-art [2], [3], [7] results. The current existing algorithms work in either of the two broadly divided classes: local translation and global translation. But neither of the algorithms excel in both photorealism and faithful style translation at the same time and for all test cases. One or the other factor gets compromised. Global stylization methods work by matching statistical factors of the pixel values [11] whereas local stylization is achieved by algorithms that find close and consistent relations between pixel values of the content and style images. Another classification is based on the algorithm's ability to translate low-level and (or) high-level features. Low-level features translation involves preservation of the intricacies in the content image while modifying the color or position with respect to the style image. Whereas high-level feature translation is the mapping of broader features which by example means day to night, summer to winter translations.

The best works proposed by Luan et al. [2] and Li et al. [7] are based on the paradigm of matching the gram matrices and makes use of semantic segmentations of the content and style images. Which take in only the content and style images as the inputs for the network. These algorithms perform post-processing like affine smoothing techniques thereby drastically improving the quality of the resultant images. Such methodologies make use of segmented images derived from the content and style images and then perform style translations from one segment to another by comparing the gram matrices of the input images. Other such algorithms based on a similar paradigm are proposed by Gatys et al. [1], Huang et al. [2] and many others, [1], [7], [12].

Promising results have been showcased by various GAN architectures namely Pix2pix by Isola et al. [13], Unsupervised Image-to-image Translation by Liu et al. [4], CycleGAN and BicycleGAN by Zhu et al. [5]. All of which take in a dataset consisting of multiple images similar to the content and the style domain. Multimodal Unsupervised Image-to-Image Translation by Huang et al. [3] provides an approach to the problem by narrowing down the content domain to only one image and a number of style images which constitute the style latent code [3]. They have proposed that to make the translation unsupervised the style images are decomposed into a common style latent space. The content space is sampled from this style space based on a conditional distribution to perform the translation.

In our proposal we narrow down our method to one content and one style image which does not make it completely unsupervised as there is only one target style image. We use the same architecture as proposed by Huang et al. [3] with an additional affine loss factor added to the loss function which adds to the smoothness and faithful style transfer which are combined with the properties of generative adversarial networks.

### III. METHODOLOGY

In addition to the model proposed by Huang et al. [3] we add the local affine transform  $\mathcal{L}_m$  also known as the photorealism factor of the content image calculated from the Matting Laplacian matrix proposed by Levin et al. [6].

#### A. Assumptions

All assumptions are exactly the same as that made in the paper Multimodal Unsupervised Image-to-Image Translation by Huang et al. [3] which are as follows. The model assumes that the content and style images are composed of distinct image spaces  $x_i \in \mathcal{X}_i$  where  $x_i$  is the  $i^{th}$  image and  $\mathcal{X}_i$  is its corresponding image space. Here our goal is to estimate the conditional distributions  $p(x_1|x_2)$  and  $p(x_2|x_1)$  leading to the learned translation models  $p(x_1 \rightarrow 2|x_2)$  and  $p(x_2 \rightarrow 1|x_1)$  respectively given that  $p(x_1)$  and  $p(x_2)$  are the marginal distributions of  $x_1$  and  $x_2$  respectively.

We make another assumption that  $x_i \in \mathcal{X}_i$  is composed of a content latent space  $c \in \mathcal{C}$  and a style latent space  $s_i \in \mathcal{S}_i$  corresponding to every image from the dataset. Thus two images  $(x_1, x_2)$  are generated from the individual generators by  $x_1 = G_1^*(c, s_2)$  and  $x_2 = G_2^*(c, s_1)$ .  $G_1^*$  and  $G_2^*$  are generator functions with  $E_1^*$  and  $E_2^*$  being their inverse encoders where  $E_1^* = (G_1^*)^{-1}$  and  $E_2^* = (G_2^*)^{-1}$ . Hence our aim is to train the encoder and generator functions using neural networks.


#### B. Matting Laplacian

Image matting is the process of extracting the foreground and the background from an image with minimal possible user intervention. The Matting Laplacian [6] process produces an alpha matte which is the segmented image with the foreground object in white and the background in black or vice versa as


per the requirements. Using this matting laplacian matrix we calculate the local affine transform factor  $\mathcal{L}_m$  also known as the photorealism factor.

$$\mathcal{L}_m = \sum_{c=1}^3 V_c[O]^T \mathcal{M}_I V_c[O] \quad (1)$$

It is a summation of the affine losses of all the three channels of the image.  $\mathcal{M}_I$  is the least-squares penalty function that is dependant on the input image  $I$ . The dimensions of the  $\mathcal{M}_I$  matrix are  $(N \times N)$  and  $V_c[O]$  is the vectorized format of the input image  $O$  in the channel  $c$  having dimensions  $(N \times I)$ . Thus this factor proves crucial in preserving the photorealism and the content image in our proposal.



(a) Single domain reconstruction of  $x_1$



(b) Single domain reconstruction of  $x_2$

Fig. 1. The above given images are the representations of how the self domain reconstruction of our model works. The images  $x_1$  and  $x_2$  are encoded into their respective content and style latent codes  $c_i$  and  $s_i$ . The reconstructed images  $x_1'$  and  $x_2'$  are not equal to their corresponding input images because of  $\mathcal{L}_1$  loss.

#### C. Model

Our model given in figure 3 constitutes an encoder and a decoder  $E_i^*$  and  $G_i^*$  respectively for every domain  $\mathcal{X}_i$ , in our case  $i = 1, 2$ . The encoder is factorized from the content and style latent codes  $c_i$  and  $s_i$ .

$$(c_i, s_i) = (E_i^c(x_i), E_i^s(x_i)) = E_i(x_i) \quad (2)$$

Thus for image-to-image translation we interchange the encoders and decoders i.e. for translation  $x_1 \rightarrow 2$  we make use of the content code  $c_1 = E_1^c(x_1)$  and a randomly drawn style latent code from  $s_2$ . Subsequently we use the decoder  $G_2$  to generate the image.

$$x_{1 \rightarrow 2} = G_2(c_1, s_2) \quad (3)$$

The loss function is composed of two factors, the bidirectional reconstruction loss and the adversarial loss. The

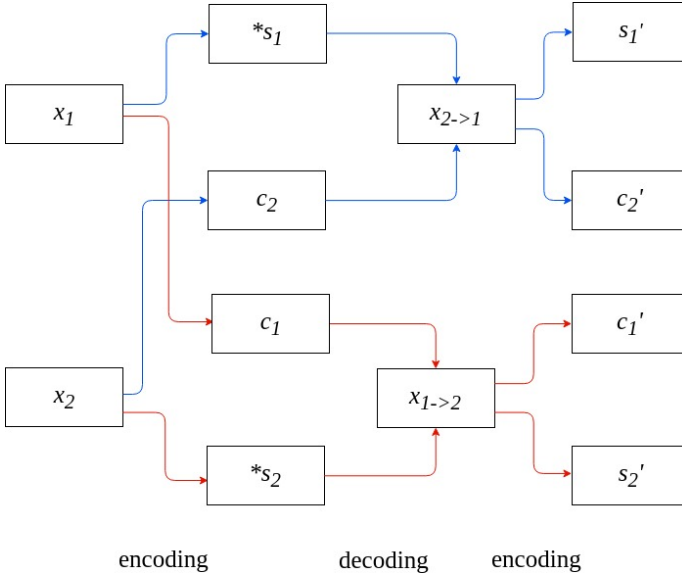


Fig. 2. This image represents the cross-domain translation of  $x_{1 \rightarrow 2}$  and  $x_{2 \rightarrow 1}$ . The '\*' represents the Gaussian prior. We encounter  $\mathcal{L}_1$  losses when reconstructing images from  $s_i$  to  $s_i'$  and  $c_i$  to  $c_i'$  thereby fulfilling the bidirectional reconstruction properties of our model. GAN loss is encountered when the translation of  $x_1 \rightarrow x_{2 \rightarrow 1}$  and  $x_2 \rightarrow x_{1 \rightarrow 2}$  takes place.

bidirectional reconstruction loss is added to make sure that there is a two way reconstruction of images in the directions, image  $\rightarrow$  latent  $\rightarrow$  image and latent  $\rightarrow$  image  $\rightarrow$  latent. The image reconstruction loss is computed as the difference between the image reconstructed from the latent spaces  $c_1$  and  $s_1$  of image  $x_1$  and the image  $x_1$  which is given by (it is similar to  $\mathcal{L}_{recon}^{x_2}$  for the image  $x_2$ ):

$$\mathcal{L}_{recon}^{x_1} = \mathbb{E}_{x_1 \sim p(x_1)} [\|G_1(E_1(x_1)) - x_1\|_1] \quad (4)$$

The latent reconstruction loss  $\mathcal{L}_{recon}^{c_1}$  is the difference between the content encoding of the generated image  $G_2(c_1, s_2)$  and the content encoding  $c_1$  of the image  $x_1$  and  $\mathcal{L}_{recon}^{s_2}$  is the difference between the style encoding of the generated image  $G_2(c_1, s_2)$  and the style encoding  $s_2$  of the image  $x_2$  they are given by the equations (which are similar for and  $\mathcal{L}_{recon}^{s_1}$ ):

$$\mathcal{L}_{recon}^{c_1} = \mathbb{E}_{c_1 \sim p(c_1), s_2 \sim p(s_2)} [\|E_2^c(G_2(c_1, s_2)) - c_1\|_1] \quad (5)$$

$$\mathcal{L}_{recon}^{s_2} = \mathbb{E}_{c_1 \sim p(c_1), s_2 \sim p(s_2)} [\|E_2^s(G_2(c_1, s_2)) - s_2\|_1] \quad (6)$$

Here  $q(s_2)$  is defined as the prior  $\mathcal{N}(0, \mathbf{I})$  and  $p(c_1)$  is defined as  $c_1 = E_1^c(x_1)$  where  $x_1 \sim p(x_1)$ .

Since we use a GAN framework we encounter an adversarial loss which is supposed to be minimised so that the generated images are as identical as possible to the original images. This loss is given by:

$$\mathcal{L}_{GAN}^{x_2} = \mathbb{E}_{c_1 \sim p(c_1), s_2 \sim p(s_2)} [\log(1 - D_2(G_2(c_1, s_2)))] + \mathbb{E}_{x_2 \sim p(x_2)} [\log D_2(x_2)] \quad (7)$$

Here  $D_2$  is the discriminator function that distinguishes between the real image  $x_2$  and the translated images. The discriminator function  $D_1$  and loss  $\mathcal{L}_{GAN}^{x_1}$  are defined in a similar way.

As mentioned earlier the architecture we use is essentially the same as that was proposed by Huang et al. [3]. The only difference being that in our approach we use only one style image and add the affine transform loss in the overall loss function. We assume [2] that the input images are photorealistic and we do not have to lose this property. Thus we penalize the loss function with the photorealism factor so as not to lose this property while minimizing the reconstruction losses from the image, content and style latent spaces. The overall loss function proposed by us is given by:

$$\min_{E_1, E_2, G_1, G_2, D_1, D_2} \max \mathcal{L}(E_1, E_2, G_1, G_2, D_1, D_2) = \mathcal{L}_{GAN}^{x_1} + \mathcal{L}_{GAN}^{x_2} + \lambda_x (\mathcal{L}_{recon}^{x_1} + \mathcal{L}_{recon}^{x_2}) + \lambda_c (\mathcal{L}_{recon}^{c_1} + \mathcal{L}_{recon}^{c_2}) + \lambda_s (\mathcal{L}_{recon}^{s_1} + \mathcal{L}_{recon}^{s_2}) + \lambda_A (\mathcal{L}_m^{x_1} + \mathcal{L}_m^{x_2}) \quad (8)$$

Where  $\lambda_x$ ,  $\lambda_c$ ,  $\lambda_s$  are the weights that control the reconstruction, and  $\lambda_A$  is the photorealism regularization weight [2].

#### D. Analysis

Our goal is to minimize the loss function defined in equation (8). This minima is the optimal state of of our model and at this point the following states are achieved:

$$p(c_1) = p(c_2) \quad (9)$$

$$p(s_1) = q(s_1) \quad (10)$$

$$p(s_2) = q(s_2) \quad (11)$$

$$p(x_1, x_{1 \rightarrow 2}) + p(c_1) = p(x_{2 \rightarrow 1}, x_2) + p(c_2) \quad (12)$$

The equation (12) is different from the one proposed by Huang et al. [3] because our model adds the local affine loss of the content images. Our model is constructed in such a way that when  $x_1$  is the content image  $x_2$  is taken as the style image and vice versa. Which is why the local affine loss of both the images is taken into consideration in equation (8) and also the content marginal distributions are added and taken into account when comparing the joint distributions  $p(x_1, x_{1 \rightarrow 2})$  and  $p(x_{2 \rightarrow 1}, x_2)$ . At this state the content marginal distributions  $p(c_1)$  and  $p(c_2)$  also become equal. Also at this optimal state the style marginal distributions  $p(s_i)$  are equal to their prior distributions  $q(s_i)$ . The fact that we use one one-to-one image mapping makes our process sound like it follows the supervised learning paradigm, but it does not. Even though we have only one image in the content and style domain the images are encoded into a content and style latent space and translated on the basis of conditional probability. Thus our method is free from any deterministic translations as performed by the methods [1], [2], [7], [12], [14] which make use of image segmentation that helps in mapping regions of interest in both the content and style images.

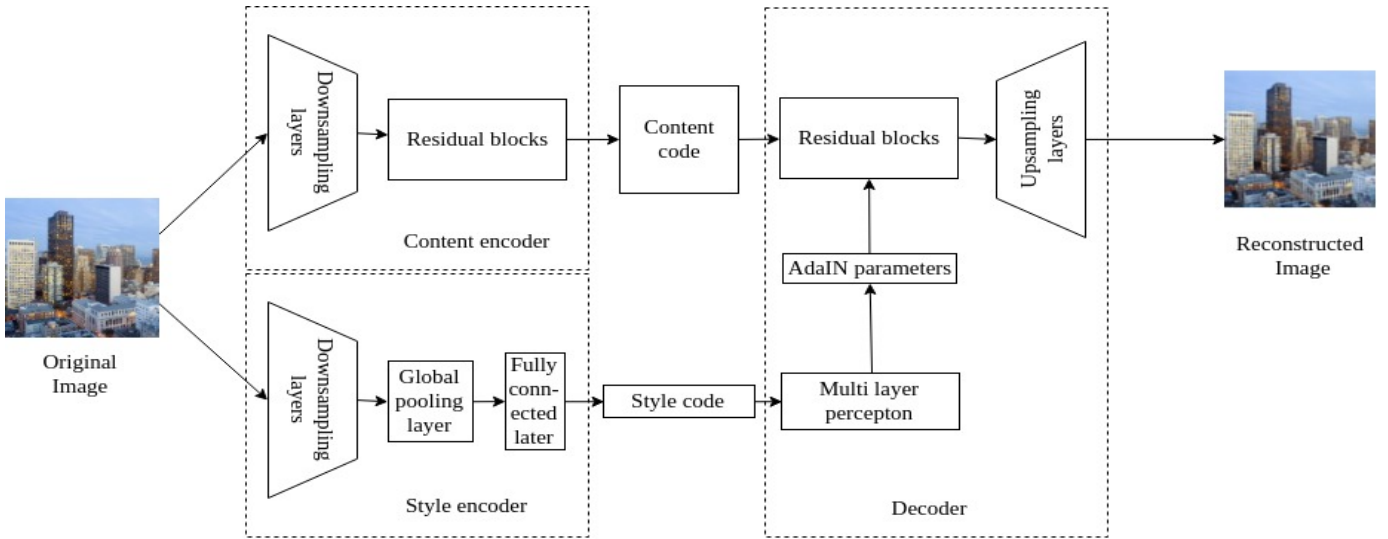


Fig. 3. The auto-encoder architecture. It consists of a content encoder, style encoder and a decoder. The content encoder comprises of three convolutional blocks which perform downsampling, followed by four residual blocks. The style encoder comprises of five convolutional layers followed by a global average pooling layer followed by a fully connected layer at the end. The decoder makes use of a multi-layer perceptron that generates AdaIN [15] parameters from the style code. The content code along with AdaIN parameters is processed by four residual blocks. The output of the residual blocks is passed on to three upsampling layers that generate the final image.




Fig. 4. First two images in each row i.e. [4a and 4b], [4g and 4h] and [4m and 4n] are the constituent datasets. Image 4c and 4i and 4o are the results obtained from Multimodal Unsupervised Image-to-image Translation [3] with 4a, 4g and 4m as the content images respectively and 4b, 4h and 4n as the style images respectively. While 4d, 4j and 4p are the results produced by the same method with 4b, 4h and 4n as the content images respectively and 4a, 4g and 4m as the style images respectively. Image 4e is our result with 4a as the content image and 4b as the style image. Image 4f is our result with 4b as the content image and 4a as the style image. Image 4k is our result with 4g as the content image and 4h as the style image. Image 4l is our result with 4h as the content image and 4g as the style image. Image 4q is our result with 4m as the content image and 4n as the style image. Image 4r is our result with 4n as the content image and 4m as the style image.

#### IV. IMPLEMENTATION DETAILS

We have adapted the publicly available pytorch implementation of Multimodal Unsupervised Image-to-image Translation [3]. The architecture consists of an auto-encoder (generator) and a discriminator. The auto-encoder comprises of a separate content and style encoder and a combined decoder. The auto-encoder architecture consists of the following layers:

- The content encoder whose content makes up the content latent space (in the listed order):
  - $7 \times 7$  convolutional block with stride 1 and 64 filters.
  - $4 \times 4$  convolutional block with stride 2 and 128 filters.
  - $4 \times 4$  convolutional block with stride 2 and 256 filters.
  - 4 residual blocks each consisting of two  $3 \times 3$  convolutional blocks with 256 filters.
- The style encoder whose output is added to the style latent space (in the listed order):
  - $7 \times 7$  convolutional block with stride 1 and 64 filters.
  - $4 \times 4$  convolutional block with stride 2 and 128 filters.
  - $3 \times 4 \times 4$  convolutional block with stride 2 and 256 filters.
  - Global average pooling layer.
  - Fully connected layer with 8 filters.
- The decoder which reconstructs an image from the content and style latent code (in the listed order):
  - 4 residual blocks each consisting of two  $3 \times 3$  convolutional blocks with 256 filters.
  - $2 \times 2$  nearest-neighbour upsampling layer followed by a  $5 \times 5$  convolutional layer with stride 1 and 128 filters.
  - $2 \times 2$  nearest-neighbour upsampling layer followed by a  $5 \times 5$  convolutional layer with stride 1 and 64 filters.
  - $7 \times 7$  convolutional block with stride 1 and 3 filters.

The discriminator used is a multi-scale discriminator proposed by Wang et al. [9] which makes use of the LSGAN objective function proposed by Mao et al. [10]. This helps to pilot the generator towards producing realistic and perform effective translation while preserving the content. The architecture consists of the following layers in the listed order:

- $4 \times 4$  convolutional block with stride 2 and 64 filters.
- $4 \times 4$  convolutional block with stride 2 and 128 filters.
- $4 \times 4$  convolutional block with stride 2 and 256 filters.
- $4 \times 4$  convolutional block with stride 2 and 512 filters.

We use the python implementation to compute the Matting Laplacian matrix [16] from the tensorflow implementation of Deep Photo Style Transfer [2]. The image, content and style reconstruction weights and the photorealism regularization weight are experimentally set to  $\lambda_x = 10$ ,  $\lambda_c = 1$ ,  $\lambda_s = 1$  and  $\lambda_A = 10^4$  [2] respectively. Our implementation is available on <https://github.com/ozamanan/semisit>.

#### V. RESULTS

Our dataset is composed of only two 3 channel images with resolution  $256 \times 256$ . Thus we use a batch size of 1. Furthermore for every iteration both the images from the dataset are used once to train the respective parts of the network. At once when one image is used as the content image the other one is used as the style image and vice versa thereby completing the bidirectional reconstruction process. All images used for experimental purposes are taken from the implementation of Deep Photo Style Transfer [2].

The images fig. 4e, 4f, 4k, 4l, 4q and 4r shown in fig. 4 are the results generated from our code whereas the images 4c, 4d, 4i, 4j, 4o and 4p are the results generated using the code of Huang et al. [3]. The results shown by us are the optimal results beyond which the images tend to converge to their respective style images. The optimal solution is the one where the resultant image holds the properties of both the content and style images while still being recognised by the discriminator as a constituent image of the dataset. This optimal state is mentioned in eq. (12).

This optimal state clearly shows an improvement in content preservation and image smoothness over the proposal of Huang et al. [3]. It is achieved due to the addition of the affine transform factors  $\mathcal{L}_m^{x_1}$  and  $\mathcal{L}_m^{x_2}$ . Thus our proposed methodology generates results that are better in comparison to the results from the method used by Huang et al. [3].

#### VI. CONCLUSION

We have proposed an architecture that performs the task of unsupervised image-to-image translation with better accuracy and results. The future work includes reducing the noise and making the results more accurate even for low resolutions. Another future scope lies in broadening this architecture for the generation of music, text and videos.

#### REFERENCES

- [1] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2414-2423, 2016.
- [2] Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. In: CVPR. (2017)
- [3] Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: ECCV. (2018)
- [4] Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: NIPS. (2017)
- [5] Almahairi, A., Rajeswar, S., Sordoni, A., Bachman, P., Courville, A.: Augmented cycleGAN: Learning many-to-many mappings from unpaired data. arXiv preprint arXiv:1802.10151 (2018)
- [6] A. Levin, D. Lischinski, and Y. Weiss. A closed-form solution to natural image matting. IEEE Transactions on Pattern Analysis and Machine Intelligence, 30(2):228-242, 2008.
- [7] Y. Li, M.-Y. Liu, X. Li, M.-H. Yang, and J. Kautz, A closed-form solution to photorealistic image stylization, arXiv:1802.06474, 2018.
- [8] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS. (2014)
- [9] Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High resolution image synthesis and semantic manipulation with conditional GANs. In: CVPR. (2018)
- [10] Mao, X., Li, Q., Xie, H., Lau, Y.R., Wang, Z., Smolley, S.P.: Least squares generative adversarial networks. In: ICCV. (2017)

- [11] Erik Reinhard, Michael Ashikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer Graphics and Applications*, 21(5):3441, 2001.
- [12] Ulyanov, V. Lebedev, A. Vedaldi, and V. Lempitsky. Texture networks: Feedforward synthesis of textures and stylized images. In *International Conference on Machine Learning (ICML)*, 2016.
- [13] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. 2016.
- [14] Paulina Hensman and Kiyoharu Aizawa. cgan-based manga colorization using a single training image. arXiv:1706.06918, 2017.
- [15] Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: *ICCV*. (2017)
- [16] Martin Benson. [Online]. Available: [https://github.com/martinbenson/deep-photo-styletransfer/blob/master/deep\\_photo.py](https://github.com/martinbenson/deep-photo-styletransfer/blob/master/deep_photo.py)