

Content and Colour Distillation for Learning Image Translations with the Spatial Profile Loss

M. Saquib Sarfraz^{1,2}
saquib.sarfraz@kit.edu

Constantin Seibold¹
constantin.seibold@kit.edu

Haroon Khalid¹
haroon.khalid49@gmail.com

Rainer Stiefelhagen¹
cvhci.anthropomatik.kit.edu

¹ Institute of Anthropomatics & Robotics
Karlsruhe Institute of Technology
Germany

² Autonomous Systems
Daimler TSS
Germany

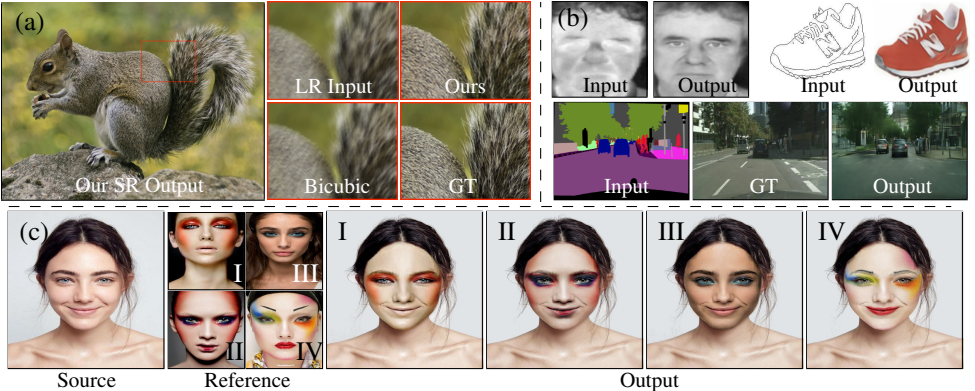


Figure 1: Various results synthesized by training single generator with our proposed loss formulation. a). Single Image super resolution b). Image-to-image domain mappings and c). Photo realistic reference makeup transfer: We can generate a HD 2048×2048 portrait with transferred makeup for a given low resolution 512×512 source and reference pair.

Abstract

Generative adversarial networks has emerged as a defacto standard for image translation problems. To successfully drive such models, one has to rely on additional networks e.g., discriminators and/or perceptual networks. Training these networks with pixel based losses alone are generally not sufficient to learn the target distribution. In this paper, we propose a novel method of computing the loss directly between the source and target images that enable proper distillation of shape/content and colour/style. We show that this is useful in typical image-to-image translations allowing us to successfully drive the generator without relying on additional networks. We demonstrate this on many difficult image translation problems such as image-to-image domain mapping, single image super-resolution and photo realistic makeup transfer. Our extensive evaluation shows the effectiveness of the proposed formulation and its ability to synthesize realistic images.

1 Introduction

Generative Adversarial Networks (GANs) by Goodfellow et al. [16] have dominated the image translation applications [11, 20, 29, 43, 46, 47]. Adversarial training is challenging as it requires two networks to compete with each other. No matter how attractive the idea of pitting two networks against each other sound, studies have demonstrated the difficulty attached to the process of adversarial training for finding the saddle point that is highly susceptible to hyperparameter settings [8, 32, 35]. A lot of recent research focus has been on stabilizing this process [6, 17, 33, 37, 38]. Since the discriminator is seen as the primary driving force, its loss is defined to minimize the distance between the training distribution and the generated distribution. Originally, Jensen-Shannon divergence was used as a distance metric [16]. A few more stable alternatives including least squares [33] and Wasserstein distance [2] are also commonly used. Primarily, the goal is to train the generator, the discriminator is only trained to steer the loss in the right direction through the generator and is discarded once the generator has been trained. Adversarial training has emerged as a default choice for image translations involving one-to-many (e.g., label-to-image) or many-to-many (image-to-image) mappings. For the generator in GANs, a pixel-based loss such as L1 is used on the output of the network [20, 26]. Some improved loss functions have been proposed and used in many recent GAN based frameworks, among which a perceptual-loss [21] (computed over features of a pretrained CNN model) have been found very effective. However, as noted in [20, 26] these losses alone are not sufficient and are used as added assistance with the adversarial training.

In this paper, we propose a new loss formulation termed *the Spatial Profile Loss* which is measured in a way that it approximates measuring the distance between the two distributions while still being based on the pixel-based comparisons. We compute this on image gradients' space and on converted colour spaces during training and show that this has a net effect of distilling shape/content from colour/style. The proposed loss is meaningful enough to directly drive the generator for one-to-many or many-to-many image translation problems. We discuss some related work in the next section. In Section 3 we describe our proposal in detail along with its applications on diverse image translation problems.

2 Related Work

Image-to-image translation: Image-to-image translation describes a broad category of problems that transfer an input image from one domain to another. In comparison to previous feed-forward approaches that used regression terms as reconstruction losses [11, 36], adversarial approaches improve the representation of frequency structure and replaced these direct feed-forward methods [22, 24, 39, 40, 41]. Isola *et al.* [20] propose a generic framework using image-conditional GANs for general image-to-image translation such as label-to-photo mappings. In their approach, they combine adversarial losses with an L1 reconstruction loss, which leads to a better representation of low-frequency structures. This basic framework was used as a stepping stone for multiple GAN-based image-to-image translation methods. For example, Zhu *et al.* [46] replace the standard L1-loss with cycle-consistency- and identity-losses by introducing another generator-discriminator pair for an unpaired setting. While adversarial approaches produce visually convincing results, the training process can be rather difficult and a balance between reconstruction and adversarial losses has to be found [32]. Furthermore, as multiple networks are included, one is more restricted regarding usable resources, which in return can negatively impact the training [8, 32]. Our work differs, as additional networks are not required, which in return opens up more possibilities

regarding the network architecture and simplifies the training process.

Recent approaches adopt perceptual losses (either reconstruction or as adversarial) by extracting an image’s high-level features of a pre-trained image classification network [9, 32, 39, 40, 41]. In doing so, the generators manage to create structure expected by a classification network. For example, Chen and Koltun [9] showed promising results on label-to-image mapping problem using cascaded refinement networks (CRN). Here, a single generator is trained using a modified perceptual loss. Similarly, Mechrez *et al.* [32] proposed a contextual loss for image transformation tasks, which represents the differences of distributions of features extracted by an external network such as a pretrained VGG19. In contrast, we propose that one can simply train a single network without the influence of such external information by investing in different loss formulations than those used at present.

Single-image super-resolution: Super-resolution has received significant attention and has been studied for decades. Super-resolution (SR) is the process of the estimation of a high-resolution image from its low-resolution counterpart. SR methods that employ deep learning have shown significant quantitative improvement in terms of PSNR values [12, 19, 23, 28]. In literature, photorealism has mostly been affiliated with GANs and now the SR problem is also being addressed with various GANs architecture in many recent works [9, 42, 44]. Recently, Ledig *et al.* [24] (SRGAN) showed state-of-the-art results on SR trained with both perceptual and adversarial losses. We make use of SRGAN’s generator and train it to restore SR images without adversarial and perceptual objectives.

Makeup Transfer and removal: Makeup transfer can be regarded as a sub-problem of style transfer, where the goal is to transfer facial makeup of a reference face image onto an input face. Liu *et al.* [30] follow the route of neural style transfer [15] and transfer the makeup locally on facial components. GANs now have produced photorealistic results on this problem. Most recently adapted CycleGAN variants [7, 25] were used for both makeup transfer and makeup removal. While Li *et al.* [25] use a symmetric approach using a segmentation level makeup loss as an additional perceptual loss, Chang *et al.* [7] approach this problem in an asymmetric manner thus using distinctive local generator-discriminator pairs for both makeup transfer and removal. We show how we can directly train a standard generator for this task as well, producing photo-realistic results without the need of such an extensive adversarial training.

3 Image Translations with the Spatial Profile Loss

For image translation problems, it is important to preserve high frequency details while transferring low frequency content (e.g., uniform textures, colour, tone etc.) as much as possible. The purpose of the generator is to produce an image as close as possible to the target. While in the typical GANs this is primarily achieved using the discriminator loss, a reconstruction term is also added at the generator output to minimize the difference between the target and generated image. The reconstruction loss is simply the pixel-based L2- or L1-loss. Since this is the average of pixel-wise differences, it is not that effective in reconstructing the actual high frequency content and the low frequency details in the image.

To illustrate this point, Figure 2 shows two very different patterns of size 64×64 (B,C), which have the same number of white pixels. Average pixel difference such as computed in L1-loss may not be of much information about the structural changes if the underlying average differences are about the same. If we compute a mean absolute difference between the source A and each of these targets, we get the same value of 0.25. Image content or structure can be characterized better by viewing the pixel variations along a given direction




Figure 2: Patterns B & C are very different but their L1 difference with A is the same.

in the image. Along this direction, the changing pixel values define a particular pattern that describes the underlying pixel relations. In our example, if we consider each image column as a vector and compute the Euclidean distance between corresponding column vectors in two images, the average of these distances comes out to be 4 between A & B and 2.8 between A & C. One can see that such a matching on *spatial profiles* can capture better structural and perceptual content in the image while being more informative.

We propose a novel loss formulation around these ideas. In principle, we compute the loss between a source and target image by considering such pattern differences along the image x and y -directions. Considering a row or a column spatial profile of an image as a vector, we can compute the similarity between them in this induced vector space. Formally, this similarity is measured over each image channel 'c' as:

$$\mathcal{L}(x, y | G(x; \theta), y) = \sum_c \left(\frac{1}{H} \text{trace}(G(x; \theta)_c \cdot y_c^\tau) + \frac{1}{W} \text{trace}(G(x; \theta)_c^\tau \cdot y_c) \right) \quad (1)$$

where $(\cdot)^\tau$ represents transpose. The first term computes similarity among row profiles and the second among column profiles of an image pair (x, y) of size $H \times W$. These image pixels profiles are L2-normalized to have a normalized cosine similarity loss. Given the proposed general scheme in Equation 1, we propose to define two loss components to drive the content/shape and style/colour of the images separately. Since an image shape is measured by its edge content, we can compute this loss with Equation 1 in the image gradients space, termed Gradient Profile (GP) loss:

$$GP(x, y) = \mathcal{L}(x, y | \nabla G(x; \theta), \nabla y) \quad (2)$$

The image gradients in each channel can easily be computed by simple 1-pixel shifted image differences from itself. We observe in our extensive ablation studies and trainings that this GP loss is quite powerful in driving gradients with healthy magnitudes resulting in fast model training in fewer epochs. We will show that it effectively distills the shape content from the colour content of the image. Similar to this, the colour content of the image can be measured e.g., by directly using Equation 1 on RGB pixels. For problems such as a one-to-many mapping of colours, RGB tends to collapse to an average colour value. This has been studied well in image colourization problems [15] where a different colour space such as LAB is used. In contrast, we observed that training the network on converted colour space images is not that effective and propose to only compute the loss in a different colour space. For this, we use the luma-chrominance YUV conversion. On this converted space the chrominance UV channels carry the colour information and computing the loss in this space transfers the colours effectively. Additionally, computing the loss in the gradient space of the chrominance channels helps the network learn stronger colour transfer and it preserves fine details driven by the colour gradients. We show on typical one-to-many mapping problems that it effectively learns the colour mappings. Our colour loss, therefore is the sum of the three terms, the RGB, YUV and ∇YUV loss. We define our Colour Profile (CP) loss as:

$$CP(x, y) = \mathcal{L}(G(x; \theta)^{RGB}, y^{RGB}) + \mathcal{L}(G(x; \theta)^{YUV}, y^{YUV}) + \mathcal{L}(\nabla G(x; \theta)^{YUV}, \nabla y^{YUV}) \quad (3)$$

The combined loss to drive the generator network, termed as Spatial Profile Loss (SPL)




Figure 3: Loss Analysis. We train separately with shape (*GP*) only and colour (*CP*) only loss to show what they drive. We display the result at evolution of epochs during training on an image from Cityscapes validation set and also show the final output with our combined *SPL* on label2image.

is the sum of *GP* and *CP* losses. Note that since Equation 1 computes a similarity, we can minimize the negative of this as:

$$\mathcal{L}_{SPL}(x, y) = \arg \min_G -(GP(x, y) + CP(x, y)) \quad (4)$$

We show in our experiments that these two terms distills the shape and colour/tone of the image and provide control in driving the two terms with different targets, e.g., we can compute the *GP* loss between output and input and the colour loss between output and target to train a generator in an auto-encoder fashion for style-transfer applications.

3.1 Image-to-Image Domain Mapping

In this section, we show the effectiveness of *SPL* for general image-to-image domain mappings. We take the same 9 block ResNet encoder-decoder generator network used in the pix2pix and CycleGAN frameworks of Zhu *et al.* [44]. We then directly train it on 256x256 images of several domain mapping problems (night2day, the thermal2visible Carl face dataset [13], edges2shoes and Cityscapes) using the proposed \mathcal{L}_{SPL} loss. We use a minibatch size of 16 with a fixed learning rate of 2×10^{-4} for 50, 50, 150 and 150 epochs respectively.

Experiments: In Figure 3, we conduct an **ablation study** in order to demonstrate the effect of *GP* and *CP* loss component and how it distills and drive the shape and colour respectively. The figure shows the results for *SPL* trained with only *GP* or only *CP* at 1, 10, 50 and 100 epochs from left to right, making it apparent that *GP* enforces the structure of the target image while *CP* influences the colour transfer. The final combined effect of *GP* and *CP* using the \mathcal{L}_{SPL} is presented at the end which shows that *CP* loss transfers the colour properties to the shape enforced by *GP* to create the synthesized image.

Figure 4 provides comparison of different loss functions on the Cityscapes label2image

	Ground Truth	L1	pix2pix [44]	CRN-loss [8]	CX-loss [54]	SPL
Pixel acc.:	0.82	0.48	0.71	0.74	0.76	0.76
Class IOU:	0.28	0.13	0.18	0.22	0.24	0.23

Figure 4: Loss Comparison: Same architecture trained on the Cityscapes label2image with different loss functions. **Top:** Averaged FCN-8-scores are provided on the full validation set.




Figure 5: SPL on various domain mapping tasks: Day-night (Left), thermal-visible (center) and edges2shoes (Right).

task with quantitative comparisons at the top and qualitative results in the bottom. We train the pix2pix generator with only L1 loss, Pix2Pix’s losses, and state-of-the-art feature-level perceptual losses i.e., CRN [9] and Contextual (CX) [54] loss formulations. The performance is compared on the commonly used averaged FCN8-scores [20, 60] on the full validation set. The perceptual feature-level losses CX and CRN are computed by using all of the feature maps of the pre-trained VGG19 network as it produces the best results for these losses. We have used the respective author’s official loss implementations. In terms of metrics, the proposed SPL formulation not only outperforms its directly comparable pixel-level L1-loss but is also on par with the adversarial (GAN) and perceptual feature level losses.

When comparing the results qualitatively in Figure 4, our SPL formulation presents a sensible distribution of colours on different classes. Pedestrians and cars are shown in differing separated colours. Similarly, the model maps reasonable outputs onto traffic signs. The feature level losses learn distinct shapes and sharper details but have a more generalized output for each class without much variability in the textures or colour details on the objects. Nevertheless, SPL has shortcomings when it comes to bringing out sharper structure on background e.g., buildings, as it defaults to a general blurry shape. This could be because the details on the buildings are not consistent and have much more variability to learn than is achievable by computing the loss over vertical or horizontal pixel profiles. The problem could be approached, for example, by exploring other spatial profile paths that may capture this information better. Overall, SPL is able to learn one-to-many mappings while only using pixel-level information and as such not needing assisting networks – discriminators and/or perceptual-networks – to train.

Figure 5 shows qualitative results of our trained SPL model on the night2day (top row) and day2night transformation (bottom row). In the middle column, we show the results of the mapping of thermal images to visible face images. It’s quite clear that our models learn an identity-preserving mapping for each paired set of thermal and visible facial images. We provide details of this training in the supplementary. The last column shows a qualitative comparison with Isola *et al.* [20] and Wang *et al.* [69] on a sample of the edges2shoes test data. In direct comparison, the same generator trained with adversarial loss [20] has problems when mapping colour on complex input structure, while SPL creates a highly plausible colour mapping as the result stays inside the boundaries of edges representing shoelaces etc.

3.2 Single Image Super Resolution

Single Image Super-Resolution (SISR) is an important image translation problem being addressed with GANs [24]. We train the generator architecture used in SRGAN [24] and EDSR [28] directly with our SPL loss on 4 times unknown downsampled LR images. For training, we take a random patch of 96x96 from the input images to produce HR resolution




Figure 6: Super-resolution results of '0801', '0816' and '0853' (DIV2K) with scale factor $\times 4$. Full resolution ground truth (left column) with output of different methods shown. SPL (ours) achieves the state-of-the-art and restore sharper details.

Method	bicubic	FSRCNN [17]	EDSR [23]	EDSR+ [24]	SRGAN+ [24]	SPL (ours)
PSNR	22.85	22.79	22.67	25.77	24.33	26.03
SSIM	0.65	0.61	0.62	0.71	0.67	0.75

Table 1: Quantitative evaluation on DIV2K track 2 dataset of the proposed SPL, in terms of PSNR and SSIM.

image upsampled by a factor of 4. Note that the generator network can take input of any arbitrary size as it is a fully convolutional network. The network is trained from scratch for 50 epochs as compared to its SRGAN counterpart which was trained for 250 epochs.

Dataset: We use the track 2 data from DIV2K [10] dataset for training. Track 2 is highly challenging as it contains downsampled LR images with unknown degradations (noise, blurring and pixel shifting). This dataset contains 800 training images and 100 validation images which cover a variety of natural scenarios i.e. animals, person images, buildings. As with previous methods, the performance is tested on the provided 100 validation images.

Results and Comparison: We compare the performance of our proposed SPL on DIV2K track 2 with state-of-the-art supervised SISR methods: FSRCNN [17], EDSR [23] and SRGAN [24]. For quantitative evaluation, we use PSNR and structural similarity SSIM. The results for FSRCNN, EDSR and SRGAN have been taken from Yi et al. [44]. They show improved results for EDSR and SRGAN by first training these models on track 1 data (bicubic downscaling) and then fine-tuning on track 2 data (EDSR+, SRGAN+). Table 1 shows the mean PSNR and SSIM values of restored validation images. It illustrates that FSRCNN and EDSR perform worse if the blur and noise are unknown in the training process. Fine-tuning with track 2 improves the results for EDSR+ and SRGAN+. Our model being trained from scratch with SPL on track 2 achieves the state-of-the-art in comparison to the considered SISR methods Figure 6 shows a more subjective analysis of chosen validation images. For every example image, the ground truth target and restored images are displayed with cropped details in the subsequent row. As can be seen, where EDSR+ manages to get comparable




Figure 7: a). Proposed makeup transfer sequence: Source and reference image are passed to a generator. Its result is added onto the reference, resulting in the final output, over which the \mathcal{L}_{SPL} loss is computed. b). Preprocessing through warping and corresponding results.

PSNR values, it restores smoother regions lacking high frequency edge information and fine structure details. In contrast, SPL has learned to distill high frequency details through the GP loss and restored sharper colour content via the CP component of the proposed \mathcal{L}_{SPL} . We include the full SR output of these methods in the supplementary material.

3.3 Photo Realistic Makeup Transfer

We demonstrate the usefulness of distilling shape from colour and the control we have in our loss formulation on makeup transfer applications. We propose a novel and effective design for training a single network for such problems, see Figure 7(a). Here, we combine the source and the reference image to a 6-channel input. Our generator creates an intermediate result, which is added to the reference to reach the final output. We use our GP loss between the generated output and the source image while the CP loss component is computed between generator output and reference image. This way the network can easily reproduce the source face with all its structure preserved while only transferring makeup from the reference and allows us to train on unpaired data. As this approach is not restricted to a specific network, we can adjust it to any problem at hand.

We use the generator architecture of CycleGAN [46] for standard and the SR generator (used in section 3.2) for $\times 4$ SR makeup transfer tasks. For our training, we warp the reference face on the source image as shown in Figure 7(b). The warping process can lead to different types of artefacts, however, our colour transfer is mostly not affected by this. Each model is trained for 10 epochs with a batch size of 10.

Proposed FCC Dataset: Existing makeup transfer databases are relatively small [8, 18] (less than 1000 images each) or were, to our knowledge at the time, not publicly available. We, therefore, constructed a new **Facial Cosmetic Content** database (FCC) for strong makeup as well as general makeup in both higher and lower resolutions. We detected face images from Youtube makeup tutorials from various regions depicting racial diversity. Images from the first and last quarter of each video are kept in order to capture faces with and without makeup. We then classify the extracted images as strongly visible makeup, general makeup, non-makeup by extracting CNN features [9]. After removing remaining noise, we re-sized the images to either 256×256 or 512×512 . This way, we gather 5389 non-makeup and 13036 makeup images out of which 5281 contain strongly visible makeup. The FCC dataset with a total of ~ 18400 images is currently the largest dataset for makeup style transfer applications. Sample images of the dataset and more details are provided in the supplementary.

Experiments: In Figure 8, we perform an **ablation study** to analyse the influence of each loss term by separately training the generator each time with one part of the described losses



Figure 8: Loss Analysis. Comparison of impact of different loss terms for makeup transfer. From left to right: Input, reference, GP , $GP+RGB$ -profile loss, $GP+RGB+YUV$ -profile loss, full \mathcal{L}_{SPL} as well as \mathcal{L}_{SPL} with reversed inputs.

added with equal weight to the total loss. Considering only GP between source and output, the generator resurfaces the underlying structure of the source image, while removing most of its colour. After including the RGB profile loss between source and reference, the generator learns to transfer the basic colour scheme of the reference image, without being impacted by portrayed facial features like eye wrinkles or freckles. However, it cannot translate visual details as it is not able to grasp brightness variations properly. The YUV-profile loss lets the generator recreate better tone transitions. By adding the YUV-gradient loss term (∇YUV), the model learns to better represent brightness transitions, thus improving the saturation and leading to a more natural transfer. Note that because of our proposed design, interestingly, the same model trained for makeup transfer also manages to remove makeup when swapping source and reference, as depicted in the last snippet of Figure 8.

Comparison with state-of-the-art: In Figure 9, we compare our results to four different previous works. Liu *et al.* [60] appears to be affected by factors such as pose and expression variations that can lead to artefacts like asymmetric eyes or overlaps between source and reference. Fišer *et al.*'s [44] approach applied to makeup transfer can show problematic tone transitions between skin regions. Liao *et al.* [47] works well on areas sharing similar features while the resulting colour transfer suffers in the case of excessive makeup regions like eyeshadows. Most recent work by Chang *et al.* [8] train a CycleGAN variant each on eye, mouth and skin regions as the makeup between these regions can strongly vary. While this approach achieves makeup of quality comparable to the reference, the generated image might not display a correct transfer from it, as it will only recreate styles already seen in training. In contrast, our network learns to preserve the source's facial attributes such as expression and pose and only transfer the reference's makeup. Since our model only learns a colour transfer between source and reference image, we are robust to styles different to our training data and so can replicate any form of makeup ranging from only affecting small areas to full-face makeups. It is sufficient to train our single network for only 10 epochs with SPL, while Chang *et al.* [8] train a multitude of generator-discriminator pairs for at least 400 epochs. In Figure 10, we compare with Chang *et al.* [8] on an extreme makeup style that was unseen during either training. For this experiment, we show the result of our SR makeup transfer producing full HD 2048×2048 images for LR 512×512 source and reference images (see also Figure 1). While the result by Chang *et al.* [8] seems reasonable, it fails to replicate fine colour details shown by the reference. SPL fully captures the makeup




Figure 9: State-of-the-art comparison between different makeup transfer methods.




Figure 10: Comparison with Chang *et al.* [10] on extreme makeup styles. We demonstrate our SR-SPL model for makeup transfer. Eye closeup: Chang *et al.* (bottom) and ours (top).

tone on the cheeks, sparkling makeup around the bottom eyelid as well as the star-shaped structures beneath it.

We include more details and the limitations of our method in the supplementary material. Our code along with the proposed FCC dataset is available online¹.

4 Conclusion

The Spatial Profile Loss formulation makes use of both gradients and chromatic channels to approximate distances in shape and color between two images. A unique property of the proposed formulation is the use of loss calculations on pixel patterns defined by horizontal or vertical profiles. The central idea is to approximate the image distribution in a low-dimensional subspace spanned by these pixel-profiles vectors. While we have shown that simple vertical/horizontal profiles suffice for approximating the target distribution one can explore more orientations. In fact we have experimented with including the diagonal profiles, nevertheless, with negligible qualitative improvements. As a regular sampling method using specific orientations, it’s possible there are attributes of an image that are not being captured. For future work, including more paths such as random walks or other types of paths would be interesting, if not as efficient to compute.

Conclusively, we have presented an effective loss formulation based on spatial profiles for general image translation problems. The proposed SPL formulation enables us to independently drive the shape and colour which, among others, is useful in typical style/colour transfer problems. On diverse image translation tasks, our results demonstrate that the proposed loss components can effectively minimize the difference between the generated and the target image distributions. This suggests that even without adversarial training or additional perceptual networks realistic transfer is possible and can even surpass its counterpart. We hope that our work may suggest seeking alternatives and more effective ways to train the generators for similar problems.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, volume 3, page 2, 2017.
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Con-*

¹<https://github.com/ssarfraz/SPL>

- ference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
 - [4] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do image degradation first. In *European Conference on Computer Vision*, 2018.
 - [5] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face and Gesture Recognition*, 2018.
 - [6] Yanshuai Cao, Gavin Weiguang Ding, Kry Yik-Chau Lui, and Ruitong Huang. Improving GAN training via binarized representation entropy (BRE) regularization. In *International Conference on Learning Representations (ICLR)*, 2018.
 - [7] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. In *2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
 - [8] Cunjian Chen, Antitza Dantcheva, and Arun Ross. Automatic facial makeup detection with application in face recognition. In *Biometrics (ICB), 2013 International Conference on*, pages 1–8. IEEE, 2013.
 - [9] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *IEEE International Conference on Computer Vision (ICCV)*, volume 1, page 3, 2017.
 - [10] Zezhou Cheng, Qingxiong Yang, and Bin Sheng. Deep colorization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 415–423, 2015.
 - [11] Yunjey Choi, Min-Je Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
 - [12] Chao Dong, Chen Change Loy, and Xiaoou Tang. Accelerating the super-resolution convolutional neural network. In *European Conference on Computer Vision*, pages 391–407. Springer, 2016.
 - [13] Virginia Espinosa-Duró, Marcos Faundez-Zanuy, and Jiří Mekyska. A new face database simultaneously acquired in visible, near-infrared and thermal spectrums. *Cognitive Computation*, 5(1): 119–135, 2013.
 - [14] Jakub Fišer, Ondřej Jamriška, David Simons, Eli Shechtman, Jingwan Lu, Paul Asente, Michal Lukáč, and Daniel Šykora. Example-based synthesis of stylized facial animations. *ACM Transactions on Graphics (TOG)*, 36(4):155, 2017.
 - [15] LA Gatys, AS Ecker, and M Bethge. A neural algorithm of artistic style. *Nature Communications*, 2015.
 - [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
 - [17] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.

- [18] Junlin Hu, Yongxin Ge, Jiwen Lu, and Xin Feng. Makeup-robust face verification. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 2342–2346. IEEE, 2013.
- [19] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5197–5206, 2015.
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976, 2017.
- [21] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016.
- [22] Levent Karacan, Zeynep Akata, Aykut Erdem, and Erkut Erdem. Learning to generate images of outdoor scenes from attributes and semantic layouts. *arXiv preprint arXiv:1612.00215*, 2016.
- [23] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016.
- [24] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, volume 2, page 4, 2017.
- [25] Tingting Li, Ruihe Qian, Chao Dong, Si Liu, Qiong Yan, Wenwu Zhu, and Liang Lin. Beautygan: Instance-level facial makeup transfer with deep generative adversarial network. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 645–653. ACM, 2018.
- [26] Yi Li, Lingxiao Song, Xiang Wu, Ran He, and Tieniu Tan. Anti-makeup: Learning a bi-level adversarial network for makeup-invariant face verification. 2018.
- [27] Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. Visual attribute transfer through deep image analogy. *ACM Trans. Graph.*, 36(4):120:1–120:15, July 2017. ISSN 0730-0301. doi: 10.1145/3072959.3073683.
- [28] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *The IEEE conference on computer vision and pattern recognition (CVPR) workshops*, volume 1, page 4, 2017.
- [29] 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.
- [30] Si Liu, Xinyu Ou, Ruihe Qian, Wei Wang, and Xiaochun Cao. Makeup like a superstar: Deep localized makeup transfer network. In *IJCAI*, 2016.
- [31] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [32] Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, and Olivier Bousquet. Are gans created equal? a large-scale study. *arXiv preprint arXiv:1711.10337*, 2017.

- [33] 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.
- [34] Roey Mechrez, Itamar Talmi, and Lih Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 768–783, 2018.
- [35] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International Conference on Machine Learning*, pages 3478–3487, 2018.
- [36] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016.
- [37] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2234–2242, 2016.
- [38] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *CVPR*, volume 2, page 5, 2017.
- [39] Chao Wang, Haiyong Zheng, Zhibin Yu, Ziqiang Zheng, Zhaorui Gu, and Bing Zheng. Discriminative region proposal adversarial networks for high-quality image-to-image translation. In *ECCV*, 2018.
- [40] Chaoyue Wang, Chang Xu, Chaohui Wang, and Dacheng Tao. Perceptual adversarial networks for image-to-image transformation. *IEEE Transactions on Image Processing*, 27(8):4066–4079, 2018.
- [41] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [42] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Chen Change Loy, Yu Qiao, and Xiaoou Tang. Esrgan: Enhanced super-resolution generative adversarial networks. In *PIRM*, 2018.
- [43] Zili Yi, Hao (Richard) Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *ICCV*, pages 2868–2876, 2017.
- [44] Yuan Yuan, Siyuan Liu, Jiawei Zhang, Yongbing Zhang, Chao Dong, and Liang Lin. Unsupervised image super-resolution using cycle-in-cycle generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018.
- [45] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European Conference on Computer Vision*, pages 649–666. Springer, 2016.
- [46] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017.
- [47] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, pages 465–476, 2017.