

Instagram Filter Removal on Fashionable Images

Furkan Kınlı¹ Barış Özcan² Furkan Kırac³
Video, Vision and Graphics Lab
Özyeğin University

{furkan.kinli¹, furkan.kirac³}@ozyegin.edu.tr, baris.ozcan.10097@ozu.edu.tr²

Abstract

Social media images are generally transformed by filtering to obtain aesthetically more pleasing appearances. However, CNNs generally fail to interpret both the image and its filtered version as the same in the visual analysis of social media images. We introduce Instagram Filter Removal Network (IFRNet) to mitigate the effects of image filters for social media analysis applications. To achieve this, we assume any filter applied to an image substantially injects a piece of additional style information to it, and we consider this problem as a reverse style transfer problem. The visual effects of filtering can be directly removed by adaptively normalizing external style information in each level of the encoder. Experiments demonstrate that IFRNet outperforms all compared methods in quantitative and qualitative comparisons, and has the ability to remove the visual effects to a great extent. Additionally, we present the filter classification performance of our proposed model, and analyze the dominant color estimation on the images unfiltered by all compared methods.

1. Introduction

In recent years, hundred of millions of photos have been shared on different social media platforms (e.g. Instagram, Facebook, etc.). Predicting the preferences of the users by analyzing their posts on such platforms becomes a crucial element of increasing the profitability for different industries (e.g. food, e-commerce and fashion). This analysis is roughly composed of understanding contents of a bunch of images and extracting particular information that may be useful for a certain domain. The first part of this composition, *visual understanding*, is mostly attacked by various deep learning approaches and their recent developments [8].

Visual understanding for a specific purpose may contain several different tasks to be accomplished, such as image classification, object detection, instance segmentation and image retrieval. Convolutional Neural Networks (CNNs) show an outstanding performance on visual un-




Figure 1: Failure cases arising from filtering on detection and segmentation tasks in fashion domain. Example images are predicted by Attr-Mask-RCNN trained on Fashionpedia dataset, which is introduced in [15].

derstanding tasks, and there are several prominent studies [9, 10, 12, 19, 27, 28, 31] proposing different fundamental solutions to these tasks by employing the variants of CNNs. However, CNNs may not deliver the same performance in real-world applications, as in the standard benchmark studies, due to the varied distractive factors like noise or blurring in real-world images or different transformations applied to the images. Recent studies [1, 11, 21, 33] have shown that CNNs are sensitive to these distractive factors, and they lead to degrade the performance significantly.

With the intent of exhibiting aesthetically more pleasing appearances or scenes, the images shared on social media platforms are mostly transformed into a different version by applying some filters. These filters modify the origi-

nal image in many different ways (*e.g.* adjusting contrast, brightness, hue, saturation; introducing different levels of blur and noise; applying color curves or vignetting). These modifications not only make an image more aesthetically pleasing, but also interpolate a particular style information to its feature maps. Considering the performance of visual understanding systems, it is important to handle such interpolations without turning the problem into an ill-posed problem. Note that detecting the levels of modifications applied during application of the filters for each single image is an ill-posed problem. At this point, CNNs do not give the exact outputs for an image and its filtered version, as shown in Figure 1, due to the differences on their feature maps caused by from filtering. This leads to the performance degradation on visual understanding tasks in various domains.

Previous studies addressing this issue have proposed different solutions that try to classify the specific filter applied to the images [2, 3, 4, 33] or to learn the complex parameters of a set of transformations applied to the images [1, 29]. Although these solutions have the ability to predict the filter or a set of transformations applied, they could not recover the original image. In this study, we assume that any filter applied to an image basically stands for the additional style information injected to the images. Therefore, we consider this problem as a *reverse* style transfer problem. In our approach, the style information from the filtered source image is learned in adaptive manner, and normalized its feature maps in the encoder to be able to generate unfiltered images with the help of adversarial learning. Our main contribution in this study can be summarized as follows:

- We propose a novel filter removal architecture, namely *Instagram Filter Removal Network (IFRNet)*, which has encoder-decoder structure that normalizes the style information in the encoder to remove the visual effects of Instagram filters.
- We introduce a new dataset, namely *Instagram-Filtered Fashionable Images (IFFI)*, which is a set of 10,200 high-resolution fashionable images composed of 600 collected images and their filtered versions by 16 different Instagram filters.
- We compare the qualitative and quantitative results of IFRNet with the previous filter removal approach [1] and the fundamental image-to-image translation studies [14, 30, 38].
- As additional experiments, we present the filter classification performance of IFRNet, and analyze the dominant color estimation on the images unfiltered by the compared methods.

2. Related Works

2.1. Filter Recognition and Removal

There are limited studies on recognizing the filters applied to an image. Chen *et al.* [3] introduces a Siamese CNN architecture to classify the filters by employing a discriminative pair sampling and an adaptive margin contrastive objective function. Bianco *et al.* [2] investigates the performance of some standard CNN architectures (*i.e.* AlexNet [20], LeNet [22] and GoogLeNet [32]) on small subset of Places2 dataset [37] where 22 Instagram filters are applied to the samples. Likewise, Chu and Fan [4] examines the performance of transfer learning approach on AlexNet [20], VGG16 [31] and ResNet-50 [10] architectures. Sen *et al.* [29] proposes a method using CNNs that learns to extract the parameters of transformations of filters from a reference image, and then transfers this information to a target image. Wu *et al.* [33] demonstrates that it is possible to reduce the effects of filtering on image classification by employing adaptive feature normalization approach. Lastly, Bianco *et al.* [1] discusses a noticeable strategy to remove the filters from the images. This strategy involves learning the parametric local transformations for each filter adaptively by CNNs to restore the images. Apart from these studies, we introduce an adversarial methodology that directly learns to remove the visual effects brought by the filters, and recover the images back to their original versions.

2.2. Style Transfer

Style Transfer is the variant of image-to-image translation tasks where the main goal is to transfer the style information extracted from a reference image into a target image while preserving its context information. Recent studies [6, 7, 13, 14, 16, 17, 38] show that *many* style information from different reference images can be successfully transferred into *many* different target images in varied domains. The underlying common strategy in these studies is to capture the style information from the feature representation of an image, and then to learn to synthesize this information and the context information of another image. Inspired from this strategy, we assume that the filters applied to images can be interpreted as the style information injected to the original version, and with the help of adaptive feature normalization [13], it can be swept away from the images during the feature extraction. Therefore, we refer this task as *reverse style transfer* where the particular style is removed from an image, instead of transferring it into a target image.

3. Methodology

Our proposed architecture, namely *IFRNet*, directly removes the style information injected to the images by filters. IFRNet achieves it by normalizing the feature maps




Figure 2: Original images from IFFI dataset and their filtered versions by 7 example filters.

of filtered images, and directly generates the unfiltered version of them. Moreover, we introduce a dataset containing 10,200 filtered high-resolution fashionable images to validate the performance of our approach.

3.1. IFFI: Instagram-Filtered Fashionable Images

Previous studies have picked a small subset of well-known datasets (*i.e.* [5, 37]) to use for validating their performance. These subsets with the filtered versions of the samples have not been open-sourced, and it may not be possible to reproduce the results in these studies. Moreover, the samples in these datasets do not seem like social media posts, which mostly have aesthetic concerns. Due to these reasons, we have collected a set of aesthetically pleasing images that are filtered by 16 Instagram filters. *IFFI* dataset contains high-resolution (1080×1080) 600 images and with their 16 different filtered versions for each. In particular, we have picked mostly-used 16 filters: *1977*, *Amaro*, *Brannan*, *Clarendon*, *Gingham*, *He-Fe*, *Hudson*, *Lo-Fi*, *Mayfair*, *Nashville*, *Perpetua*, *Sutro*, *Toaster*, *Valencia*, *Willow*, *X-Pro II*. Some examples of collected images and their filtered versions can be seen in Figure 2.

3.2. IFRNet: Instagram Filter Removal Network

We define the problem of removing Instagram filters from the images as a reverse style transfer problem, where any visual effect injected by a particular filter is removed from an image by directly reverting them back to its original style. To achieve this, we propose *Instagram Filter Removal Network (IFRNet)*, which has encoder-decoder structure employing adaptive feature normalization strategy to

all layers in the encoder.

Style extractor module contains five-layer fully-connected network f_{fc} . It maps the feature representations \mathbf{z}_{vgg} , encoded by a pre-trained VGG network, to the latent space. There are N fully-connected heads attached to f_{fc} , where N equals to the number of normalization layer in the encoder. Each head is responsible for adapting the affine parameters y_i of corresponding normalization layers in the encoder. This module can be formulated as follows

$$y_i = h_i(f_{fc}(\mathbf{z}_{vgg})) \quad (1)$$

where $h_i(\cdot)$ is i^{th} fully-connected head of style extractor module and y_i represents the predicted mean and variance vectors for the corresponding normalization layer.

Adaptive instance normalization (AdaIN) [13] makes it possible to adapt the style of an input image to any arbitrary style of a target image by transferring the feature statistics computed across spatial locations. Formally, AdaIN receives the feature maps of the content image x and the style input y , and aligns the channel-wise mean and variance of these maps to match the style information obtained from y .

$$\text{AdaIN}(x, y) = \sigma(y) \left(\frac{x - \mu(x)}{\sigma(x)} \right) + \mu(y) \quad (2)$$

In our design, we consider the filters applied to the images as an external style information, and simply employ AdaIN to eliminate this information from these images. The encoder of IFRNet is composed of 6 residual blocks, and each of them has specific AdaIN layer to normalize the feature maps in each level with its corresponding affine parameters extracted by previous module. To preserve any

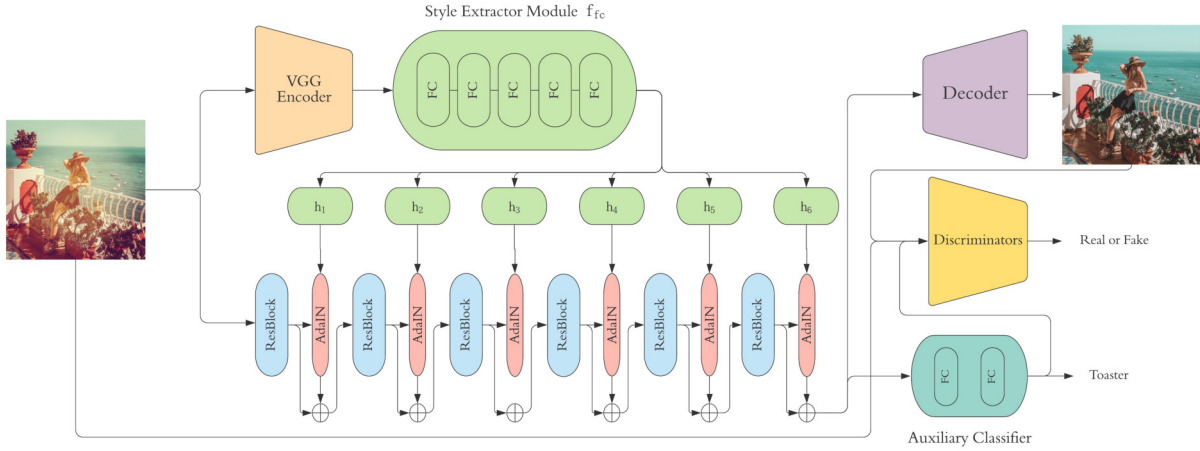


Figure 3: Overall architecture of IFRNet. The latent representations of the images are fed into style extractor module f_{fc} , then the affine parameters are calculated, and sent to the corresponding AdaIN layer to eliminate the style information from the feature maps. At the end, the decoder part of IFRNet generates the unfiltered version of the input images.

related information about the original style, we include skip connections to the normalized feature maps before sending them to the next layer. In this way, the encoder has the capacity not only to learn to remove the external style information, but also to preserve the related ones.

$$o_i = r_i(v_i, y_i) + v_i \quad (3)$$

where $r_i(\cdot)$ represents the residual block at i^{th} layer of the encoder, which receives the feature maps v_i and the affine parameters y_i , and o_i is the output of this block. Note that the normalization does not have any impact on the feature maps when its corresponding affine parameters are set to zero for a particular layer, which means there is no external style information remaining on feature maps at this level.

After obtaining the latent representations of the filtered images including no external style information, we feed them to both the decoder part of IFRNet and an auxiliary classifier for filter type. The decoder contains a number of consecutive upsampling and residual convolutional blocks, and it generates the unfiltered version of the input image with the help of adversarial training. Moreover, the auxiliary classifier is a simple feed-forward network that produces the predicted filter type. At this point, we design two discriminator models (*i.e.* PatchGAN) which separately penalize the global image and the local structures at the scale of patches, and follow the adversarial training strategy in [14]. Overall architecture of IFRNet is shown in Figure 3.

The objective function for IFRNet is composed of three main components, which are texture consistency loss, semantic consistency loss, adversarial loss extended with auxiliary classification loss. Texture consistency loss computes

the sum of the patch-wise relative difference between the feature maps of the input and the target (*i.e.* ID-MRF loss) [24, 25], and leads to enhance the details in the generated images by minimizing the inconsistency between the input patch and the most similar target patch. Secondly, semantic consistency loss [16] ensures the input and target images have similar feature representations at different scales by minimizing Euclidean distance between them (see Eq. 4).

$$\mathcal{L}_{sem} = \sum_{p=0}^{P-1} \frac{1}{C_i H_i W_i} \|\Phi_i^p(\mathbf{I}_{out}) - \Phi_i^p(\mathbf{I}_{gt})\|_2^2 \quad (4)$$

where $\Phi^p(\cdot)$ is the feature map of p^{th} layer of pre-trained VGG16 network [31] with the shape of $C_i \times H_i \times W_i$, \mathbf{I}_{out} is the output of IFRNet and \mathbf{I}_{gt} is the input image. At this point, we use only `relu3_2` layer of VGG16 for this loss with the intent of controlling over the computational complexity of training. Moreover, IFRNet contains an auxiliary classifier that receives the latent representations of the filtered images as the input, and predicts their filter types. These predictions provide the extra information to the discriminators about the input images, and thus it makes training of discriminators more stabilized. Note that employing an auxiliary classifier does not improve the performance significantly. The general formula of our adversarial training can be seen in Equation 5.

$$\mathcal{L}_{adv} = \mathcal{L}_{glo} + \mathcal{L}_{loc} + \lambda_{gp} \mathcal{L}_{gp} + \lambda_{cls} \mathcal{L}_{cls} \quad (5)$$

where \mathcal{L}_{glo} and \mathcal{L}_{loc} represent the objective functions for global and local discriminators, respectively. \mathcal{L}_{gp} is the gradient penalty whose weight λ_{gp} is set to 10, and \mathcal{L}_{cls} is the




Figure 4: Comparison of the qualitative results of Instagram filter removal on IFFI dataset. Filters applied (top to bottom): *Sutro*, *Willow*, *Nashville*, *Amaro*, *Lo-Fi*, *Toaster*.

classification loss whose weight λ_{cls} is set to 0.5. Finally, our objective function for IFRNet can be represented as follows:

$$\mathcal{L} = \lambda_{tex}\mathcal{L}_{tex} + \lambda_{sem}\mathcal{L}_{sem} + \lambda_{adv}\mathcal{L}_{adv} \quad (6)$$

3.3. Experimental Setup

In our study, we used our dataset, namely IFFI dataset, which contains 500 training and 100 test images combined with the set of their filtered versions with 16 different Instagram filters. We trained IFRNet with the resized images (*i.e.* 256×256) for 120,000 steps with the batch size of 8. We only applied horizontal flipping to the images, no other

data augmentation technique is applied. We used Adam optimizer [18] with $\beta_1 = 0.5$ and $\beta_2 = 0.9$ for both generator and discriminator networks. The learning rate for generator network is 2×10^{-4} , and 10^{-3} for discriminator networks, and we did not change it during training. The weights of the main components of our objective function can be seen as follows: $\lambda_{tex} = 10^{-3}$, $\lambda_{sem} = 10^{-4}$ and $\lambda_{adv} = 10^{-3}$. The implementation of IFRNet is done on PyTorch [26]. The experiments have been conducted on 2x NVIDIA RTX 2080Ti GPUs, and a single run takes approximately 2 days to be completed. The source code can be found at <https://github.com/birdortyedi/instagram-filter-removal-pytorch>.

4. Results

In this study, we compare the performance of IFRNet against a recent approach using CNNs for removing Instagram filters [1], and also well-known paired/unpaired image-to-image translation studies including pix2pix [14], CycleGAN [38] and AngularGAN [30]. For all these different methods, we have conducted several experiments on IFFI dataset with their default training settings.

4.1. Qualitative Comparison

As shown in Figure 4, our proposed IFRNet achieves superior performance on Instagram filter removal when compared to the other methods. IFRNet can remove the visual effects injected to the images by filters at a large scale, while the other methods produce the outputs with significant differences from their original versions (*i.e.* inconsistent background or foreground colors, some artifacts on the corners and the residuals of the filters). Therefore, we demonstrate that the external style information can be swept away from the images by normalizing the feature maps at each level in the encoder.

It is important to remark that some filters (*e.g.* *Toaster*, *Sutro*, *Willow*, *Brannan*) apply several different transformations to the images, which may substantially alter the important details in the images. For example, *Toaster* adds vignette and burning effects to the image (see the last row in Figure 4), or *Willow* directly attacks to the color information, and transform the image into a purplish gray-like image (see the second row in Figure 4). Although the compared methods struggle to recover the images filtered by such challenging filters, IFRNet is able to remove these filters within a certain extent. Note that IFRNet does not have an objective function specialized to colorize the images, as in [36, 34], hence it cannot recover the images filtered by *Willow* as well as the other filters. More examples recovered by IFRNet are shown in Figure 7.

4.2. Quantitative Analysis

We have employed four common image similarity metrics in our experiments to evaluate the quantitative performance of IFRNet, they can be specified as SSIM, PSNR, Learned Perceptual Image Patch Similarity (LPIPS) [35] and CIE 2000 Color Difference (CIE- ΔE) [23].

Table 1 summarizes the performances of our proposed model and the other compared methods trained on IFFI dataset. Our proposed model outperforms the other compared methods on all common metrics. Particularly, the previous filter removal approach using CNNs [1] has limited performance, especially on the metrics prioritizing to measure the structural and perceptual similarity between the images (*i.e.* SSIM and LPIPS), since it does not take any advantage of adversarial training. *CycleGAN* [38] suffers from its inherent design, which is suitable for unpaired

Method	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	CIE- ΔE \downarrow
PE [1]	0.748	23.41	0.069	39.55
pix2pix [14]	0.825	26.35	0.048	30.32
CycleGAN [38]	0.819	22.94	0.065	36.59
AngularGAN [30]	0.846	26.30	0.048	31.11
IFRNet (ours)	0.864	30.46	0.025	20.72

Table 1: Quantitative performance of our proposed model and other compared methods on IFFI dataset.

many-to-many image translation tasks. Although this approach shows outstanding performance on numerous applications in different domains, it falls behind our approach, and also *AngularGAN* [30] (its following study specialized for color constancy) and even *pix2pix* [14] (its primary study proposing a general solution for the image-to-image translation tasks). This intrinsically indicates that the approach learning to translate the input image to the target domain with the help of cycle consistency loss may not work as well as the other methods in the case of the existence of the paired images for many-to-one translation tasks. Note that we have used the official implementations of the compared studies with their default hyper-parameter settings.

4.3. Filter Classification

In addition to removing Instagram filters from the images, we also demonstrate the performance of the auxiliary classifier of IFRNet on the same experimental setup. As we discuss in Section 3.2, IFRNet has an auxiliary classifier that provides the extra information from the images to the discriminators to stabilize the adversarial training. This branch basically gives an output for the filter types of the images.

IFRNet achieves 87.5% overall classification accuracy on IFFI dataset. Considering classifying each filter as a separate task, we have measured the classification performance of IFRNet on each filter in order to evaluate how hard to dif-




Figure 5: Confusion matrix for the number of the correct predictions of auxiliary classifier of IFRNet for each filter on IFFI dataset.

ferentiate these filters. As shown in Figure 5, our proposed model does not perform the same for all filters, and gives better results for the filters containing different remarkable transformations (e.g. *Sutro*, *Willow* and *Toaster*). As aforementioned in Section 4.1, these particular filters dramatically change the visual details in the images, and thus it makes the process of classification of these filters less complicated. On the contrary, IFRNet gets more confused on differentiating more perceptibly similar filters, like *Amaro-Valencia*, *He-Fe-Lo-Fi* and *Clarendon-Original*. Therefore, we can say that the filters containing several different transformations are more likely to be classified correctly, but it is harder to remove them from the images.

4.4. Dominant Color Estimation Analysis

Considering the practical applications of Instagram filter removal for visual understanding tasks, we analyze the performance of a simple dominant color estimation algorithm on recovered images by IFRNet and the other methods. The algorithm applies K-Means Clustering to the color space of the images, and finds N cluster centers (i.e. the dominant colors). For our proposed model and the other methods, we compare CIE 2000 Color Difference (CIE- ΔE) [23] of all dominant colors extracted by this algorithm.

Example results of dominant color estimation algorithm on unfiltered images are shown in Figure 6. Supporting the quantitative analysis in Section 4.2, our proposed model has the ability to recover the color information better than other compared methods, **even if it does not contain any special strategy for this information**. Note that IFRNet does not directly model the color conversion mapping, but learns the injected style information. Moreover, to compare the other methods, the images unfiltered by the previous study using CNNs by exploiting the polynomial expansion [1] and AngularGAN [30] have closer and consistent dominant color predictions than the ones unfiltered by the prominent image-to-image translation methods [14, 38], since [1, 30] contain different modules to specialize to gather color information.

5. Conclusion

In this study, we introduce IFRNet, an encoder-decoder structure applying adaptive feature normalization to all levels in the encoder to remove the external visual effects injected by filters. Experiments on Instagram filter removal task verify that IFRNet eliminates the external visual effects to a great extent. The main idea behind IFRNet is to consider the external visual effects as the style information, and we show that it is possible to remove them by normalizing the style of all feature maps, similar to the approach in style transfer studies. By extending the scope of the dataset for all available Instagram filters, this method could be employed for pre-processing the social media images before feeding them into a vision framework to enhance its performance.



Figure 6: Example results of dominant color estimation on the images unfiltered by the compared methods. The bar on the right side of the images represents the weights of 5 dominant colors extracted by K-Means Clustering, and color distances (CIE 2000 [23]) to its corresponding color in the original image are shown on each color of this bar.




Figure 7: More examples from the results of IFRNet on IFFI dataset. Rows: (Odd) Filtered images, (Even) Original images at the first column, the rest represents the recovered versions of corresponding filtered images. The filters applied (left to right) for the first image: *Brannan, Lo-Fi, Perpetua, Nashville, Sutro, Valencia*, for the second image: *Brannan, 1977, Clarendon, Valencia, Gingham, He-Fe*, for third image: *Clarendon, Valencia, Brannan, Willow, Lo-Fi, Nashville*, for fourth image: *He-Fe, Brannan, Toaster, Lo-Fi, Nashville, Perpetua*.

References

- [1] Simone Bianco, Claudio Cusano, Flavio Piccoli, and Raimondo Schettini. Artistic photo filter removal using convolutional neural networks. *Journal of Electronic Imaging*, 27:1, 12 2017.
- [2] Simone Bianco, Claudio Cusano, and Raimondo Schettini. Artistic photo filtering recognition using cnns. In *International Workshop on Computational Color Imaging*, volume 10213 of *Lecture Notes in Computer Science*, pages 249–258. Springer, 2017.
- [3] Yu-Hsiu Chen, Ting-Hsuan Chao, Sheng-Yi Bai, Yen-Liang Lin, Wen-Chin Chen, and Winston H. Hsu. Filter-invariant image classification on social media photos. *MM '15*, page 855–858, New York, NY, USA, 2015. Association for Computing Machinery.
- [4] Wei-Ta Chu and Yu-Tzu Fan. Photo filter classification and filter recommendation without much manual labeling. In *2019 IEEE 21st International Workshop on Multimedia Signal Processing (MMSp)*, pages 1–6. IEEE, 2019.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [6] L. A. Gatys, A. S. Ecker, and M. Bethge. A neural algorithm of artistic style. *arXiv*, Aug 2015.
- [7] Golnaz Ghiasi, Honglak Lee, Manjunath Kudlur, Vincent Dumoulin, and Jonathon Shlens. Exploring the structure of a real-time, arbitrary neural artistic stylization network. 2017.
- [8] Yanming Guo, Y. Liu, A. Oerlemans, Songyang Lao, Song Wu, and M. Lew. Deep learning for visual understanding: A review. *Neurocomputing*, 187:27–48, 2016.
- [9] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- [11] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.
- [12] Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In Aasa Feragen, Marcello Pelillo, and Marco Loog, editors, *Similarity-Based Pattern Recognition*, pages 84–92, Cham, 2015. Springer International Publishing.
- [13] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017.
- [14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
- [15] Menglin Jia, Mengyun Shi, Mikhail Sirotenko, Yin Cui, Bharath Hariharan, Claire Cardie, Hartwig Adam, and Serge Belongie. Fashionpedia: Ontology, segmentation, and an attribute localization dataset. In *ECCV*, 2020.
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, 2016.
- [17] L. Karacan, Z. Akata, A. Erdem, and E. Erdem. Manipulating attributes of natural scenes via hallucination. *ACM Transactions on Graphics*, 39(1), Nov. 2019.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [19] Gregory R. Koch. Siamese neural networks for one-shot image recognition. 2015.
- [20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 1097–1105. Curran Associates, Inc., 2012.
- [21] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *ICLR Workshop*, 2017.
- [22] Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.
- [23] Ming Luo, Guihua Cui, and B. Rigg. The development of the cie 2000 colour-difference formula: Ciede2000. *Color Research & Application*, 26:340 – 350, 10 2001.
- [24] Roey Mechrez, Itamar Talmi, Firas Shama, and Lihi Zelnik-Manor. Learning to maintain natural image statistics, [arxiv](https://arxiv.org/abs/1803.04626). *arXiv preprint arXiv:1803.04626*, 2018.
- [25] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. *arXiv preprint arXiv:1803.02077*, 2018.
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [27] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28, pages 91–99. Curran Associates, Inc., 2015.

- [29] Mrinmoy Sen and Prasenjit Chakraborty. A deep convolutional neural network based approach to extract and apply photographic transformations. In Neeta Nain, Santosh Kumar Vipparthi, and Balasubramanian Raman, editors, *Computer Vision and Image Processing*, pages 155–162, Singapore, 2020. Springer Singapore.
- [30] Oleksii Sidorov. Conditional gans for multi-illuminant color constancy: Revolution or yet another approach? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [31] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- [32] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [33] Zhe Wu, Zuxuan Wu, Bharat Singh, and Larry Davis. Recognizing instagram filtered images with feature de-stylization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):12418–12425, Apr. 2020.
- [34] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016.
- [35] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [36] Richard Zhang, Jun-Yan Zhu, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. Real-time user-guided image colorization with learned deep priors. *ACM Transactions on Graphics (TOG)*, 9(4), 2017.
- [37] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [38] 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*, 2017.