Photorealistic Style Transfer for Automated Image Editing

Esin YAVUZ and Sébastien BÉHURET
Cyanapse Ltd., Brighton, UK – www.cyanapse.com
Correspondence: eyavuz@cyanapse.com

Abstract

Style transfer is one of the hottest topics in deep learning research. Since the original Neural Style Transfer method was published [1], various methods have been developed for transforming images from one domain to another. Among those methods, Generative Adversarial Networks (GAN) [2] took a lot of attention for their success in domain translation.

We developed a GAN-based model and a training strategy to achieve photorealistic style transfer and enable automatic image editing. With this method, a new style can be trained in less than a day on a single NVIDIA Titan Xp GPU, and inference of a single high resolution image is obtained in less than a second. The model only needs images from the two domains to learn a new style.

Our method is used in an image editing platform that will allow end users to transform their images in just a few clicks. Here we present the methodology and some examples of daytime transformations. We also use the same method for correcting imperfections in the images, such as JPEG noise, blur, haze and exposure errors, or to enhance the content of images by reducing fog and improving lighting. Using NVIDIA GPUs enabled us to build a prototype quickly and test our models in a reasonable time, with an overall speed-up factor of 14.3 compared to using a high-end CPU alone.

Motivation

Recent work on style transfer has been mostly performed on limited object domains, using transformations that usually result in discontinuities between object borders. However, photorealistic translations from one natural image domain to another requires strict reliability to the object content, and the objects that can be encountered can be more various than what is generally found in context-specific datasets. In order to develop photorealistic filters for automated image transformations, we used GANs with a number of improvements to enforce picture quality. We trained the network on our own image dataset that includes scenes and objects that are likely to be found in casual photography, and used data augmentation.

Benchmarks

Training time (batch of 4 images)	Patches of 256x256	**Inference time** (1 image)	High Resolution (2400x2400)	360° (2048x4096)
NVIDIA Titan Xp	1.4 sec	NVIDIA Titan Xp	0.84 sec	1.22 sec
INTEL Xeon E5-2697 v4	20.0 sec	INTEL Xeon E5-2697 v4	14.2 sec	20.5 sec

Methodology

GAN Architecture for style transfer

The network learns to transform images in domain A into images in domain B. It is similar to CycleGAN architecture [3], with additional loss terms to ensure that original image features are conserved, and initial training to enable learning of a variety of natural image textures.

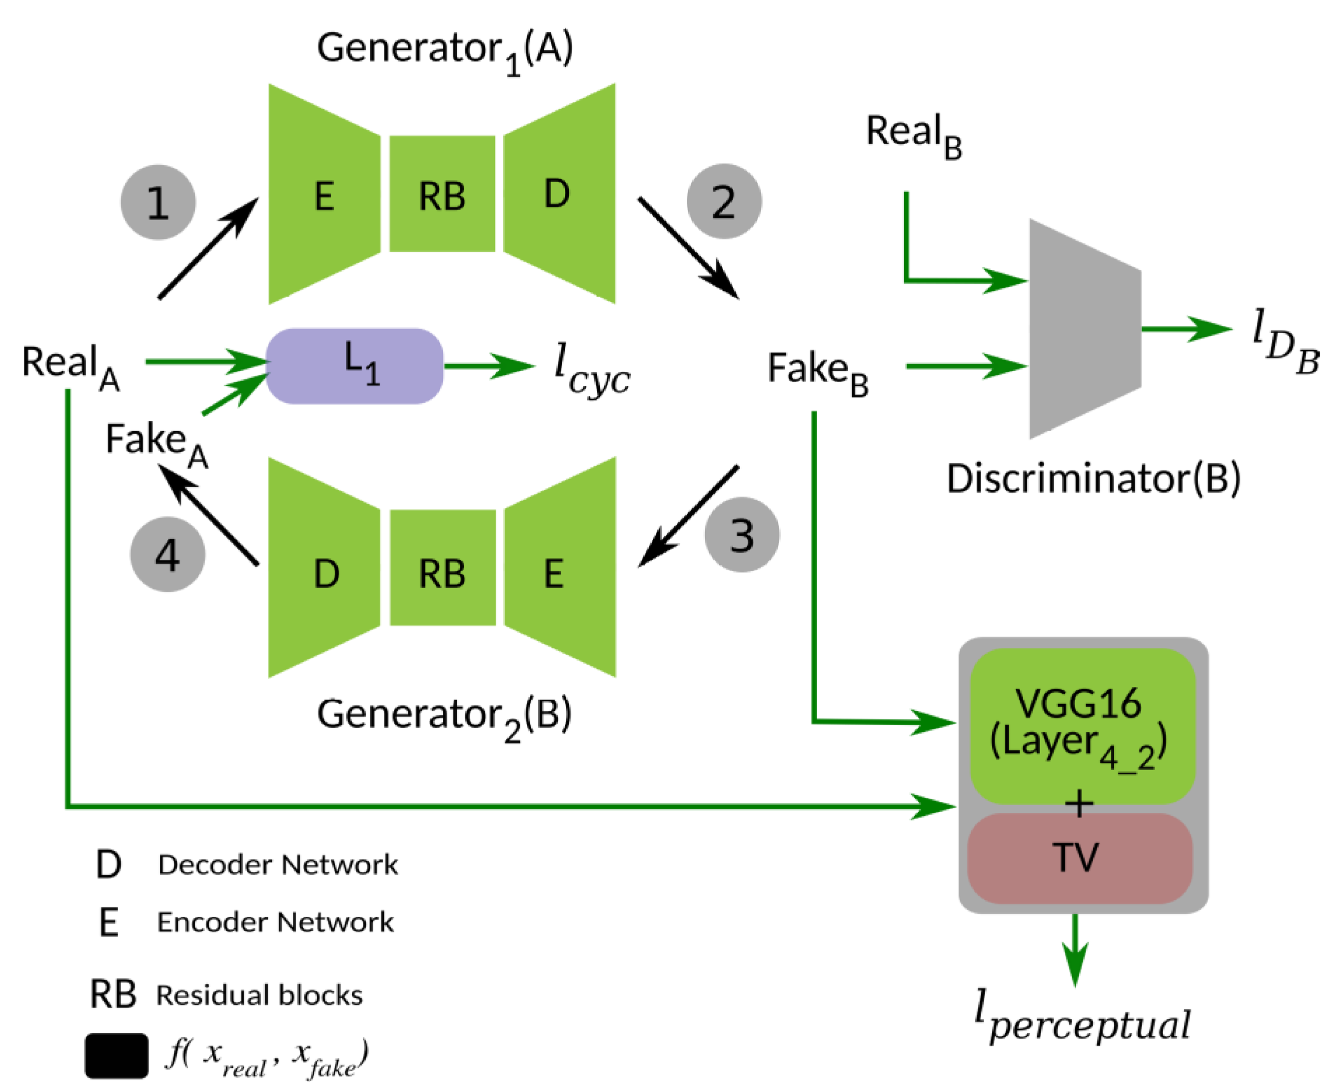

Additional improvements, while trading execution time

Warm restarts of the learning rate [4] not only helps stabilizing the learning, but also reduces artifacts. With warm restarts, we average multiple inferences obtained with different sets of weights, where each set corresponds to the end of a training cycle that has converged to a different local minima.

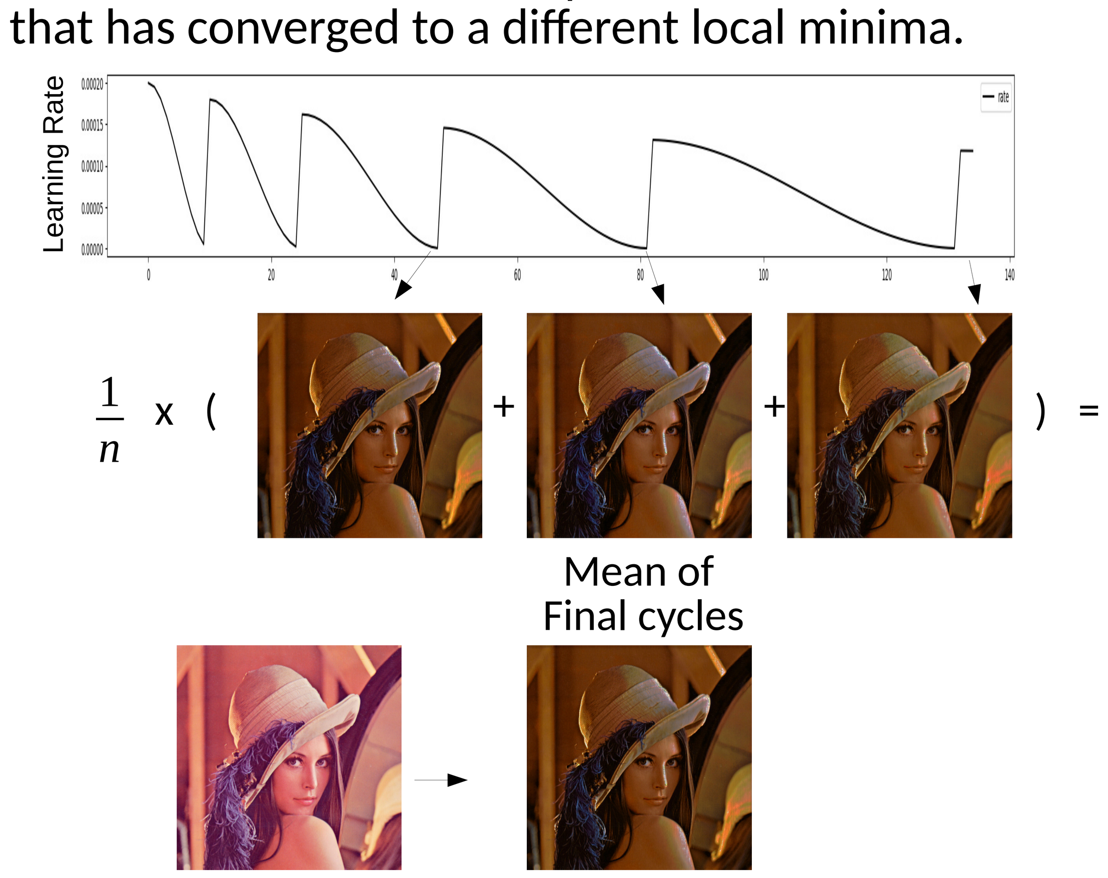

Examples

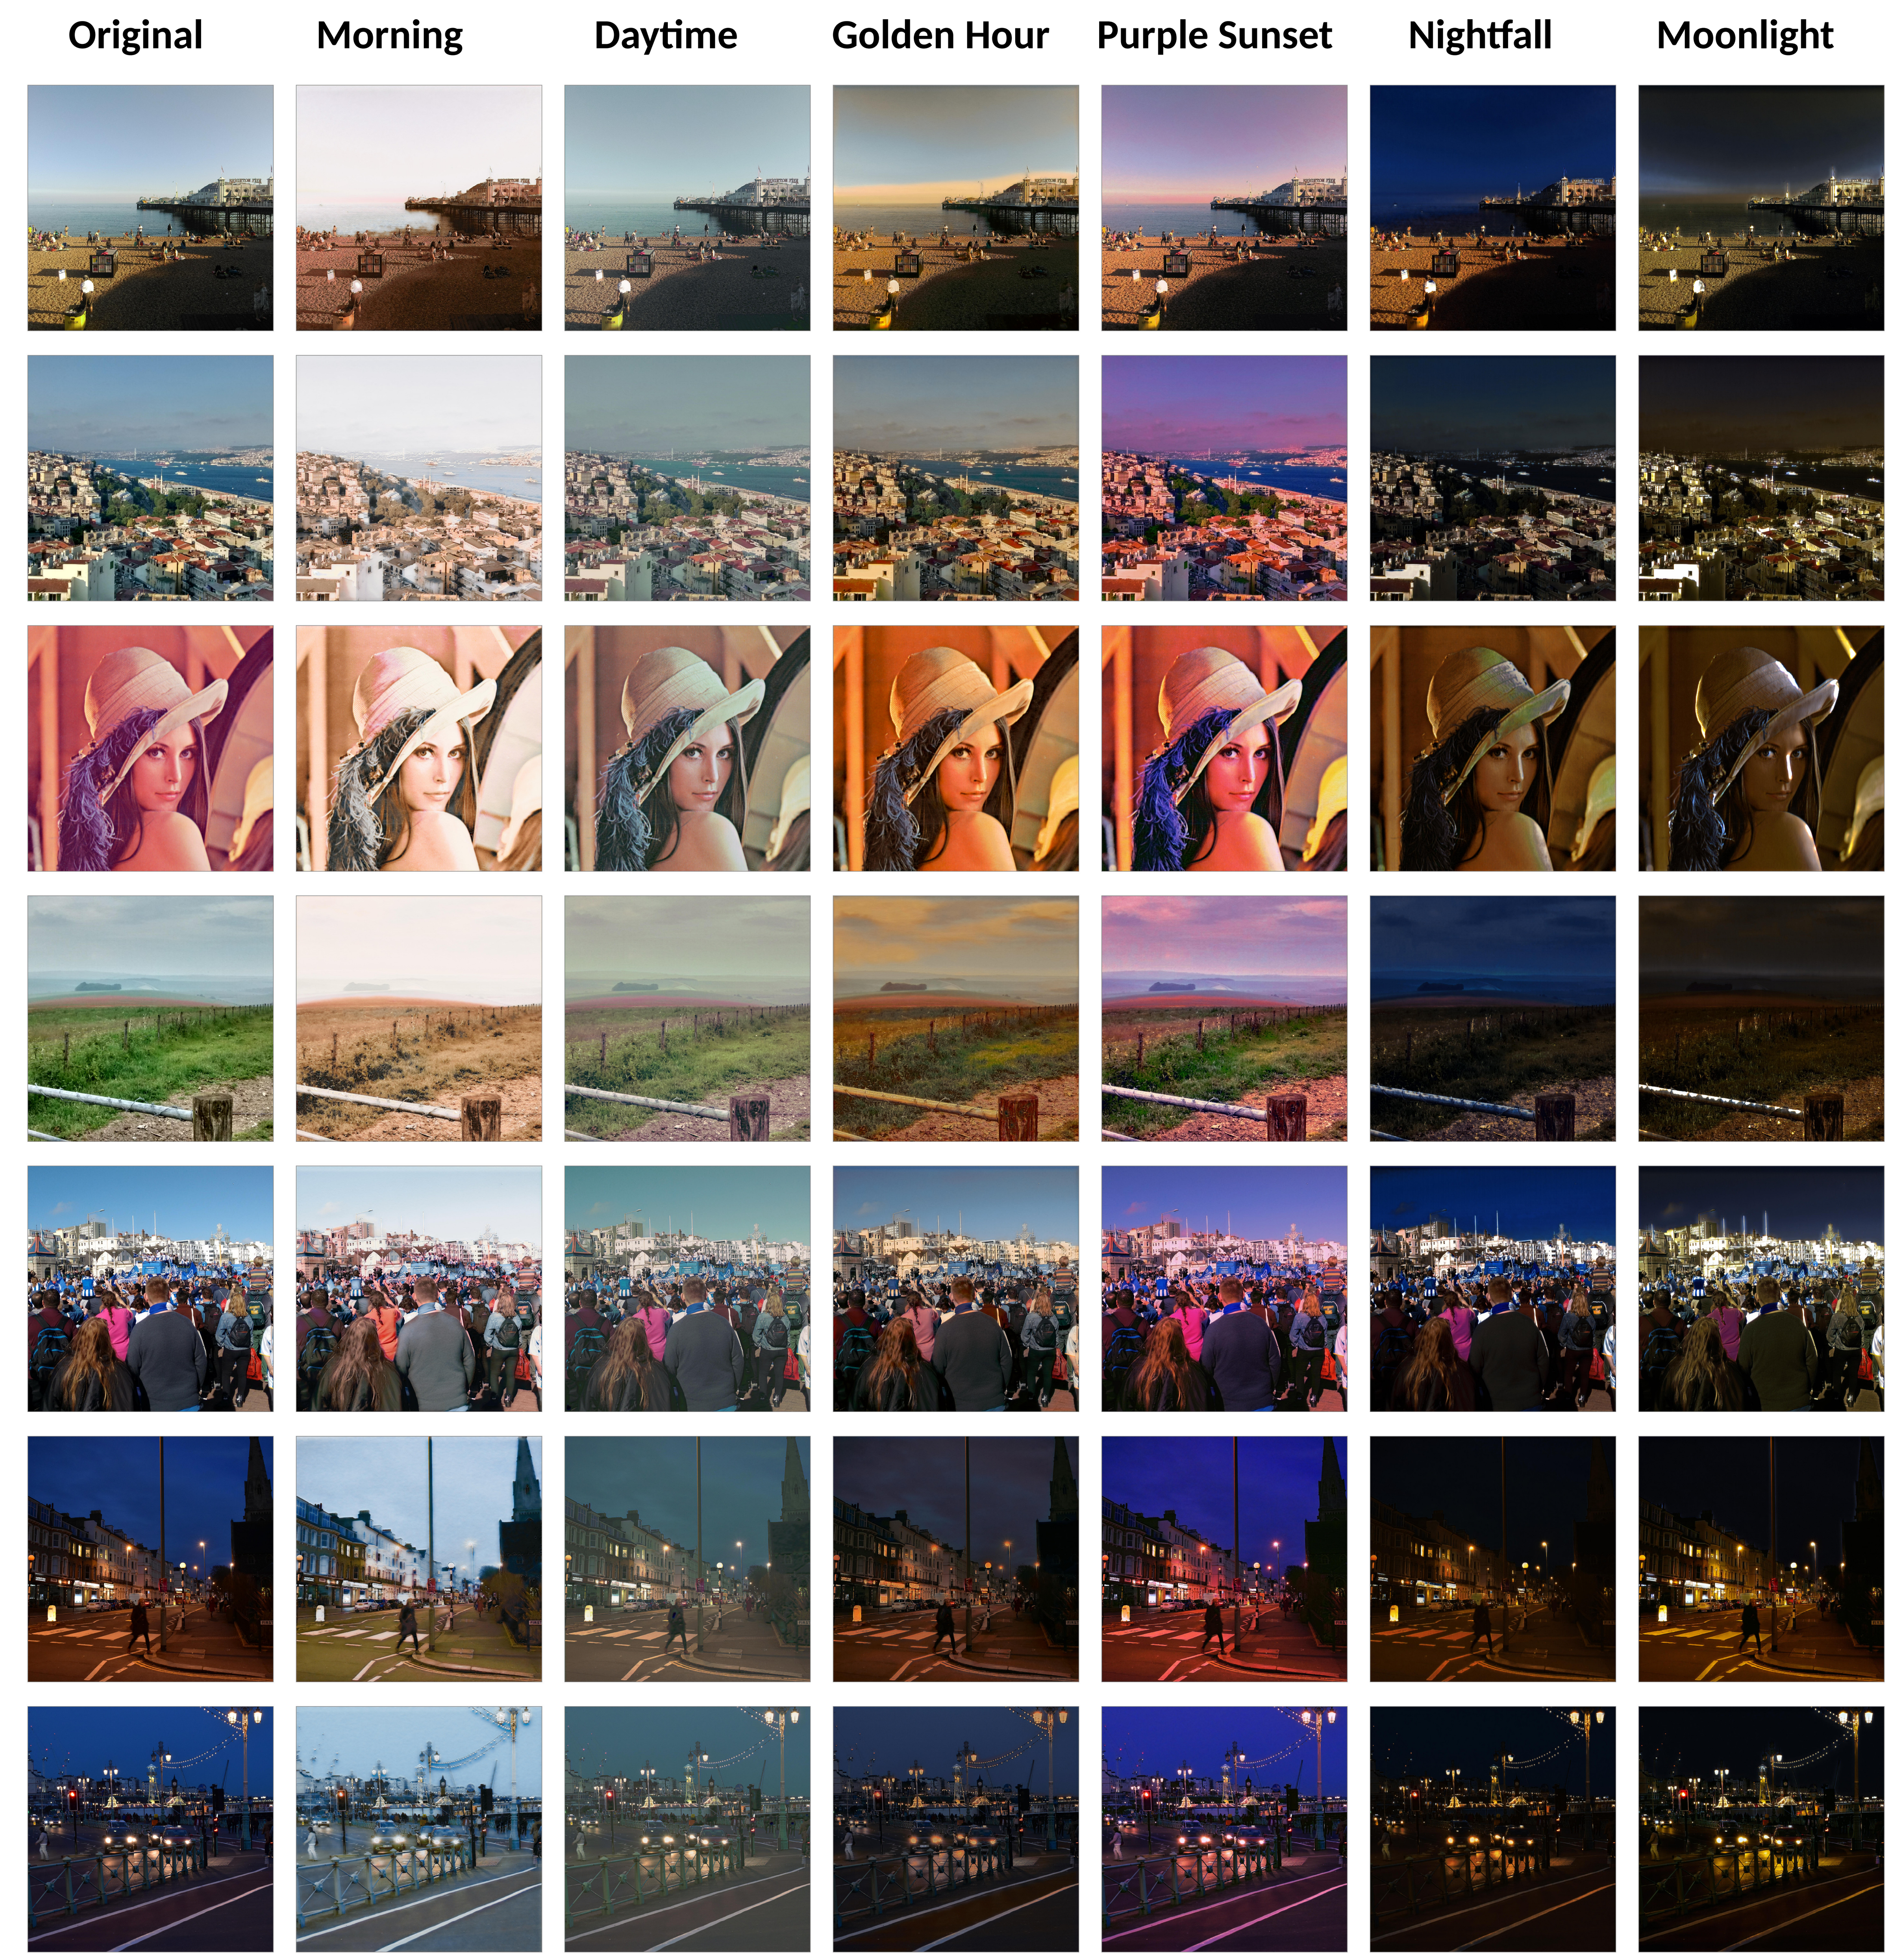

Visit www.cyanapse.com/style-transfer for full resolution images and additional examples.

Summary

Image editing tasks are likely to be fully automated in the next decade as a result of the recent advances in AI. Our platform allow fully automated editing tasks such as daytime transformation, dehazing and exposure correction in just a second using NVIDIA GPUs, without losing the original content. The platform will also be accessible through an API for bulk editing.

References

1. Gatys, L.A., Ecker, A.S. and Bethge, M., 2016. Image style transfer using convolutional neural networks. In Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on (pp. 2414-2423). IEEE.
2. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y., 2014. Generative adversarial nets. In Advances in neural information processing systems (pp. 2672-2680).
3. Zhu, J.Y., Park, T., Isola, P. and Efros, A.A., 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. arXiv preprint arXiv:1703.10593.
4. Loshchilov, I. and Hutter, F., 2017. Fixing Weight Decay Regularization in Adam. arXiv preprint arXiv:1711.05101.