WEP Warwick Evans Publishing

Transactions on Computer Science and Intelligent Systems Research
ISSN: 2960-1800, eISSN: 2960-2238 | Volume 5
AIDML 2024

Progress and Impediments in Deep Learning-Driven Image Style Transfer

Yuqi Jiang *

School of Electrical and Electronic Engineering, Chongqing University of Technology, Chongqing, China

* Corresponding Author Email: 1136881165@stu.cqut.edu.cn

Abstract. This detailed review explores key developments and ongoing challenges in image style transfer, emphasizing the transformative role of deep learning approaches, notably Convolutional Neural Networks (CNNs) and Generative Adversarial Networks (GANs). The paper examines the fundamental principles of this field, particularly the intricate process of blending 'style' and 'content' via advanced neural network designs. It chronicles critical breakthroughs, and progresses to contemporary solutions addressing real-time execution, diversity enhancement, and stability in results. Emerging techniques such as Deep Feature Interpolation and Multi-Scale Style Transfer are also scrutinized, offering insights into potential research directions. The review not only traces the technical evolution but also considers the wider impact of image style transfer, underscoring its significance in bridging art and technology. This intersection is demonstrated through applications that span from digital art creation to innovative adaptations in medical imaging.

Keywords: Image Style Transfer, Deep Learning Perspective.

1. Introduction

In the dynamic realm of deep learning, image style transfer has become a pivotal area within computer vision research. This technique, which involves transferring stylistic elements from one image to another, allows the creation of images with novel aesthetic attributes [1]. Its applications span a wide range, including art, design, and entertainment, opening new avenues for creative and innovative expression [2, 3]. Moreover, its significance extends to scientific research, offering insights into human visual perception and contributing to the field of neuroscience [4,5]. Although the initial exploration into image style transfer began later domestically than on the international stage, there has been rapid progress recently. Domestic researchers have developed various deep learning-based algorithms for image style transfer, presenting their work at renowned international conferences and in prestigious journals.

Internationally, research in this area commenced earlier and has cultivated a more developed and comprehensive research ecosystem [6-8]. Prominent research institutions and universities across the globe have contributed significantly, fostering the technology's advancement and widespread integration. Future developments in image style transfer are anticipated to focus on improving efficiency, stability, and realism [9-11]. The advent of innovative deep learning models, like Generative Adversarial Networks (GANs), is expected to further propel advancements in this sector [12]. This paper endeavors to offer a succinct yet thorough overview of the image style transfer technology, amalgamating and scrutinizing research from both domestic and international spheres. The aim is to delineate the technology's potential applications and challenges across various sectors, thereby serving as an invaluable reference for scholars and practitioners in related disciplines. It includes an in-depth discussion of the core principles and methodologies underlying image style transfer, an overview of pertinent deep learning algorithms, examples of practical applications in fields like art, design, and entertainment, and an investigation into the future trajectories and challenges within this technological domain [13, 14].

2. Basic Concepts and Principles of Image Style Transfer

The concept of image style transfer is underpinned by the distinct definitions of 'style' and 'content' within an image, along with the goals driving the transfer of style. In academic terms, 'style' is characterized as the ensemble of visual elements inherent in an image, encompassing aspects like texture, coloration, and brushwork, often reflecting an artistic expression. 'Content', in contrast, refers to the discernible components within the image, such as objects, scenes, and figures, which together articulate the central theme and meaning of the image. The core aim of style transfer technology is to apply the stylistic elements of one image onto the content structure of another, creating an output that exemplifies a new aesthetic [15].

The methodology of style transfer encompasses three stages: Initially, a deep convolutional neural network extracts distinct feature representations, encapsulating both style and content attributes. Subsequently, a synthesis phase merges these distinct style and content features, giving birth to an image with a fresh aesthetic. The final phase involves developing an optimal loss function to measure the deviation between the resulting image and the set stylistic and content standards, followed by iterative optimization to refine the image towards the target aesthetic [16,17].

In the current academic sphere, several methods have been instrumental in advancing image style transfer. The neural network-based style transfer, pioneered by Gatys et al., stands out for its approach of minimizing differences in both content and style. To overcome the high computational demands and slow processing speeds of this method, a novel algorithm for fast style transfer was introduced. This approach, employing perceptual loss functions and transformation networks, markedly accelerates the style transfer process [18]. Furthermore, the advent of an arbitrary style transfer technique, incorporating a new style encoding mechanism, allows for the conversion of any chosen style image into specific style parameters, vastly expanding the scope of style transfer applications [19]. The evolution of methods to impart artistic styles onto photographic images has significantly enhanced the versatility of style transfer, showing promising potential in areas like photography and graphic design.

3. Common Image Style Transfer Methods

Dominant strategies in image style transfer mainly divide into two types: methods that employ Convolutional Neural Networks and techniques based on Generative Adversarial Networks (GANs), alongside several innovative methodologies.

3.1. Convolutional Neural Network-Based Methods

Recent advancements in image style transfer techniques, particularly those involving Convolutional Neural Networks -based methods, have significantly impacted the computational artistry landscape [20]. Unlike traditional methods that often involve significant manual labor and artistic expertise, Approaches grounded in CNNs harness deep learning techniques to integrate the content of one image with the stylistic elements of another, effectively narrowing the divide between artistic expression and technological innovation [21].

3.1.1. Introduction to GAN-Based Methods for Image Style Transfer

Convolutional Neural Networks facilitate the extraction of image features through multiple layers of convolution operations, subsequently employing these features to execute the style transfer. Within the ambit of style transfer tasks, CNNs typically process the content and style images as separate inputs, amalgamating the style features into the content image via optimization algorithms [22,23]. The strength of approaches based on CNNs is found in their proficient extraction of both fundamental and advanced image attributes, facilitating the style adaptation across varied artistic expressions. Moreover, these methods boast computational efficiency, making them suitable for style transfer tasks on large-scale image datasets. However, CNNs exhibit relative deficiencies in handling complex textures and details. Additionally, due to the constraints of optimization algorithms, the style transfer process may converge to local optima, leading to instability in the quality of resultant images [24]. CNN-based methods find extensive application across various domains, including artistic style transfer, photo enhancement, and video stylization [25, 26].

3.1.2. Uses and Consequences

A leading method in this area is Neural Style Transfer (NST), a deep learning-based technique that creates new visuals by merging the subject matter of one photograph with the artistic flair of a different one. This technique utilizes CNNs to dissect and amalgamate content and style elements from distinct images [27]. NST has led to the creation of a variety of artistic images, enabling even those without traditional artistic training to create compelling artwork. One significant development in this area has been the introduction of improved style transfer models that address issues related to color schemes, stroke strength, and image contrast adjustments. Models like the enhanced Universal Style Transfer (UST) approach, which incorporates techniques for image blending and color improvement, provide an advanced post-processing structure, thereby tackling prevalent issues in the creation of neural-based stylized artworks. Such improvements contribute to a wider range of visual effects in stylized imagery and are validated across different NST approaches [28].

Moreover, the application of CNNs extends beyond artistic creation to practical applications in different fields. For instance, in medical imaging, CNN-based style transfer methods can aid in enhancing the visual clarity of medical scans, facilitating better diagnosis and treatment planning. In facial recognition technologies, these methods can improve system robustness by generating facial images under diverse conditions [29]. In the entertainment industry, particularly in beauty camera apps and short video platforms like TikTok, local image style transfer has gained popularity, offered users novel and engaged visual experiences [30].

Despite their growing utility, CNN-based image style transfer methods face challenges such as ensuring the generated images maintain high fidelity to the original content while effectively reflecting the desired artistic style. Furthermore, achieving a balance between computational efficiency and output quality remains a critical concern, especially for applications requiring real-time processing.

As the field progresses, future research directions may include developing more efficient and adaptable style transfer algorithms, enhancing the quality of generated images, and exploring new applications in areas such as augmented reality and virtual reality. Moreover, efforts to improve model interpretability and user control over the style transfer process will likely enhance the applicability and user experience of CNN-based image style transfer techniques.

3.2. Generative Adversarial Network-Based Methods

In image style transfer, the deployment of GANs is highly regarded for their adeptness in capturing and implementing complex artistic styles, from the detailed brushstrokes of traditional paintings to the abstract qualities of contemporary art. They are capable of accommodating a broad spectrum of styles and subjects, rendering them invaluable assets for artists and designers. The implementation of GANs in this domain has spurred the creation of novel applications, allowing users to automatically and effectively infuse their photographs and images with a variety of high-quality artistic styles.

3.2.1. Introduction to GAN-Based Methods for Image Style Transfer

Generative Adversarial Networks consist of two integral components: a generator and a discriminator, which engage in adversarial training to enable the generator to produce fake data closely mirroring the real data distribution. In style transfer tasks, GANs employ the generator to amalgamate content and style images to fabricate new images, while the discriminator evaluates the quality of these generated images. The advantage of GAN-based approaches lies in their capacity to generate high-quality style transfer outcomes, particularly excelling in managing complex textures and details. Furthermore, these methods possess a potent generative capability, capable of producing variant images in numerous styles. Nonetheless, the training process for GANs is considerably intricate, necessitating substantial computational resources and time. Additionally, issues such as mode collapse in GANs may result in a lack of diversity in the generated outcomes. GAN-based methods are particularly suited to applications in artistic creation, image editing, and virtual reality.

3.2.2. Applications and Implications

In the realm of image style transfer, GAN-based methods have marked significant advancements, presenting innovative solutions that extend beyond traditional approaches. Notable among these is the Circular LBP Prior-Based Enhanced GAN, which incorporates circular local binary patterns (LBP) within the GAN generator. This method is distinct for its ability to refine the detailed textures of generated images, ensuring that the style transfers are not just superficial overlays but integrate deep, textural nuances of the style reference. The method is further enhanced by integrating dense connection residual blocks and attention mechanisms, which bolster the extraction of high-frequency features, thereby improving the overall quality of the generated images.

Another key development is the Cyclic Image Translation Generative Adversarial Network (CIT-GAN), which has proven particularly effective in multi-domain style transfer, including challenging areas like iris presentation attack detection. The framework introduces a Styling Network that learns the distinctive style characteristics of each domain represented in the training dataset. This approach not only fosters the generation of synthetic images that are true to the reference domain's style but also enhances the model's adaptability and robustness across varied domains. Additionally, the realm of cross-modal facial image synthesis has seen remarkable progress through the deployment of collaborative bidirectional style transfer networks. This method stands out for its ability to manage large modality gaps, ensuring that the synthesis retains the core content while adeptly swapping the appearance styles between different modalities. This is particularly groundbreaking in contexts where the accurate representation of facial features across different imaging techniques is crucial.

The applications of these GAN-based image style transfer methods are diverse and impactful. In medical imaging, for example, they have been pivotal in enhancing the consistency and quality of images across different equipment and modalities, which is essential for accurate diagnosis and treatment planning. In the field of facial recognition, these methods have significantly improved system robustness by generating facial images under various conditions, thus mitigating the challenges posed by variations in lighting, posture, and expressions. Moreover, the art and design industry has embraced these technological advancements, leveraging them to create new artworks that fuse the subject matter from one picture with the artistic elements of another, thereby opening up new avenues for creativity and artistic expression. This melding of technology and art has not only expanded the toolkit available to artists and designers but has also democratized the ability to create complex, stylized works of art.

In summary, GAN-based methods in image style transfer have transcended traditional boundaries, offering sophisticated solutions that cater to a wide range of applications from medical imaging to artistic creation. These advancements not only underscore the versatility and power of generative adversarial networks but also hint at the vast potential for future innovations in this field.

3.3. Innovative Approaches

In the subsequent sections, we will explore three distinct methodologies that have significantly advanced the academic field of image style transfer. These are: Deep Feature Interpolation, Attention-Based Methods, and Multi-Scale Style Transfer. Each method contributes uniquely to the ongoing research and development in the computational understanding and application of artistic styles onto source content images.

3.3.1. Deep Feature Interpolation

The process known as Deep Feature Interpolation (DFI) modifies mid-level feature representations within a neural framework to perform style adaptation. Utilizing straightforward linear interpolation among the feature matrices from both content and style sources in the deeper sections of a convolutional neural setup, DFI presents a swifter option by merging in-depth features from dual images. This approach enables the merger of content and style aspects across various levels, thereby aiding in the generation of stylistically enhanced images while maintaining the integrity of the original content's architecture. Despite its encouraging outcomes, this technique is most effective when the interpolated features are in sync with the neural network's ingrained representations.

3.3.2. Attention-Based Methods

Attention-based methods in image style transfer leverage the concept of attention mechanisms, which have been highly successful in fields like natural language processing and are now being adapted for use in image processing. These methods focus on transferring style features in a content-aware manner, ensuring that the stylistic elements from the style image are applied appropriately to the relevant parts of the content image. By using attention maps, these methods can better preserve content structure while applying style, leading to more coherent and aesthetically pleasing results. They can capture long-range dependencies within images, allowing for more detailed and contextually relevant style transfers.

3.3.3. Multi-Scale Style Transfer

Multi-Scale Style Transfer approaches tackle the image style transfer problem by operating at different scales or resolutions. By breaking down the image into several scales, these methods can apply style transfer more effectively, ensuring that both global structures and local details are accurately captured and transferred. This approach allows for a more nuanced application of style, where large-scale patterns and small-scale textures can be handled separately. Multi-scale methods typically involve processing an image at various resolutions and combining the outputs, which helps in preserving the content's fidelity while introducing stylistic elements from the style image at appropriate levels of detail.

Each of these methods brings a unique perspective to the challenge of image style transfer, leveraging different aspects of image representation and processing to achieve impressive results. Deep Feature Interpolation offers speed and simplicity, Attention-Based Methods provide content-aware stylization, and Multi-Scale Style Transfer ensures a thorough application of style across different image scales.

4. Challenges and Future Directions

The main challenges presently faced in the field of image style transfer are divided into four key areas.

4.1. Balance of Style and Content

Within the realm of image style transfer, the objective is to infuse the style from source images (for example, a painting) into destination content images (such as a photograph), all while maintaining the content image's structural integrity and identifiable characteristics. The core challenge here involves striking an optimal balance between two often conflicting goals.

4.1.1. Maintaining Content Integrity

The altered image is expected to preserve the structure, key attributes, and overall visual aspect of the original content. Typically, this is managed by reducing a content loss metric that quantifies the disparity between the content attributes of the modified image and those of the original content image, as discerned by a deep neural network.

4.1.2. Replicating Style

The altered image should effectively replicate the style of the source image, encompassing elements like textures, color schemes, and painting strokes. This usually involves minimizing a style loss metric that computes the discrepancies among the stylistic features of the transformed image and those of the model style image, a process usually based on the analysis of correlations across various layers within a neural network, as initially established by Gatys and colleagues.

4.1.3. Approaches to Addressing the Equilibrium between Style and Content

Finding the right equilibrium between style and content during the execution of style adaptation in imagery involves a nuanced approach where the preservation of content needs to be balanced against the emulation of style. This delicate balance is typically modulated through a set of hyperparameters within the style transfer algorithm. Direct control over this balance can be exerted by adjusting the weights attributed to content and style losses within the overall loss function, influencing whether the final image aligns more closely with the original content or the intended style.

Furthermore, the selection of specific neural network layers for extracting content and style features plays a critical role; deeper layers are generally responsible for capturing more abstract content details, whereas shallower layers tend to focus on stylistic aspects such as textures. The process is also significantly shaped by the optimization strategy employed, which encompasses factors like the number of iterations, learning rate, and the initial state of the image undergoing transformation. Strategies such as initiating the process with the content image tend to favor content fidelity. Additionally, the application of regularization methods like total variation loss aids in preserving the structural integrity of the transferred image, thereby indirectly fostering a more balanced integration of style and content elements. Emerging techniques, including deep feature reshaping and adaptive instance normalization, have introduced innovative ways to fine-tune this balance, enabling more refined and customizable image adaptations.

Ultimately, achieving an ideal balance between style and content requires iterative experimentation with these variables, guided by foundational and evolving methodologies within the field as illuminated by the pioneering research of Gatys et al., among others.

4.2. Real-time Processing

The primary obstacles in achieving real-time style transfer stem from the computational demands and resource needs of the involved algorithms. Initial techniques, such as those introduced by Gatys et al. in their foundational works, depend substantially on iterative optimization, which is both resource-intensive and time-consuming, hence hindering instantaneous applications.

4.2.1. Efficiency in Computation

For real-time execution of style transfer, the operation needs to finish in just milliseconds. Methods developed by Johnson and his team along with Ulyanov and his colleagues, employing straight feed-forward architectures, have made substantial strides in achieving this aim. These tactics require a neural network to be pre-trained specifically for applying styles, which allows for significantly faster image modification compared to conventional iterative techniques after the training phase is finalized.

4.2.2. Constraints Due to Hardware

The efficiency of real-time style adaptation algorithms largely hinges on the computational hardware. Advanced GPUs can drastically decrease the time needed for processing, thus enabling faster style adaptations. Conversely, achieving such speed on lower-end devices, like smartphones or tablets, poses considerable challenges.

4.2.3. Optimizing the Model

Implementing strategies such as network pruning, quantization, and knowledge distillation can streamline neural networks and accelerate analysis time, thereby enhancing the pace of style transfer.

4.3. Diversity

4.3.1. Challenges to Achieving Diversity

The process is marked by underdetermination, indicating that a given combination of content and style may be interpreted through various stylistic representations. Consequently, traditional approaches often converge on a singular solution, thereby neglecting other potential interpretations.

Furthermore, the use of pre-trained neural networks introduces inherent biases, as these networks are predisposed to favor styles and features prevalent in their initial training datasets, potentially limiting the diversity of generated images. A significant obstacle also lies in enabling users to effectively influence the style transfer process without requiring an in-depth understanding of the underlying technology or complex interface navigation. Additionally, the development of robust evaluation metrics to assess the diversity and quality of style transfer remains a challenge. Current metrics often fall short in capturing the nuanced aspects of artistic styles.

4.3.2. Approaches to Enhancing Diversity

To address these challenges, several innovative strategies have been implemented. Feature Perturbation introduces randomness into the deep feature maps of images, producing diverse outputs while preserving the integrity of the chosen style. This approach highlights the importance of balancing creativity and fidelity in style transfer. Noise Encoding is another method that employs multilevel noise encoding to generate varied outcomes. This technique modifies different levels of a neural network's representations, demonstrating the potential of nuanced manipulations in achieving diversity.

Data Augmentation has also been instrumental, especially in fields like autonomous driving. It uses style transfer to enhance model robustness under various environmental conditions. This application underscores the significance of adapting models to real-world variability. Additionally, the use of Adversarial and Variational Techniques has proven effective in encouraging the creation of unique and diverse outputs. These techniques penalize redundancy, promoting a richer array of stylized images.

4.4. Stability

Stability in the realm of image style transfer is fundamentally concerned with achieving consistent results amid varying input images or minor changes to the input. The absence of this stability can result in significant alterations in stylized outputs, triggered by slight variations in either the input image or the style reference. This inconsistency is generally undesirable. Central to overcoming this challenge is the distinct separation of content and style within neural networks. Failure in this separation can cause the original content to become unduly influenced by the imposed style, thereby compromising the desired stability. Another significant challenge in this field is the susceptibility of neural networks to adversarial threats. Such threats, through minor and often imperceptible adjustments to the input, can lead to disproportionately large changes in the output. This vulnerability underscores the need for enhanced robustness in these models.

The sensitivity of model parameters, particularly those governing the balance between content and style within the loss function, also plays a pivotal role in the final output. The primary challenge lies in calibrating these parameters to consistently yield visually pleasing results across a variety of content and styles. Moreover, extending the scope of style transfer methods to a diverse range of styles and contents remains a formidable task. Certain methods may excel in specific scenarios but fall short in others, highlighting the need for more versatile approaches. In response to these challenges, substantial advancements have been made following the initial conceptualization of neural style transfer. These developments include more robust algorithms that enhance the stability and efficacy of the style transfer processes. By refining the separation of content and style, and mitigating the effects of adversarial influences and parameter fluctuations, these advancements significantly contribute to achieving stable and uniform style transfers across an extensive array of styles and contents.

5. Conclusion

This review emphasizes the transformative role of deep learning in the realm of image style transfer, with particular focus on the utilization of Convolutional Neural Networks and Generative Adversarial Networks. These technological advancements symbolize a fusion of artistic creativity and computational power, tackling key challenges in the field. These challenges encompass striking a balance between preserving original content integrity and infusing new styles, accelerating processing speeds, enhancing the variety of outputs, and improving the reliability of image transformation procedures. In essence, the dynamic field of image style transfer, driven by the intersection of artistic innovation and computational advances, points towards a future rich with sophisticated tools for artists and technologists alike. Ongoing research in this domain is expected to yield more refined, user-focused, and versatile style transfer tools, broadening the horizons for both artistic expression and practical applications in diverse sectors.

References

[1] Gatys, L. A., Ecker, A. S., & Bethge, M. (2016). A neural algorithm of artistic style. Journal of Vision, 16(12), 326.

[2] Johnson, J., Alahi, A., & Fei-Fei, L. (2016). Perceptual losses for real-time style transfer and super-resolution. European Conference on Computer Vision (ECCV).

[3] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. Advances in Neural Information Processing Systems (NIPS).

[4] Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. IEEE International Conference on Computer Vision (ICCV).

[5] Jing, Y., Yang, Y., Feng, Z., Ye, J., Yu, Y., & Song, M. (2020). Neural style transfer: A review. IEEE Transactions on Visualization and Computer Graphics, 26(11), 3365-3385.

[6] Luan, F., Paris, S., Shechtman, E., & Bala, K. (2017). Deep photo style transfer. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

[7] Ulyanov, D., Vedaldi, A., & Lempitsky, V. (2016). Instance normalization: The missing ingredient for fast stylization. ArXiv preprint arXiv: 1607.08022.

[8] Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR).

[9] Li, C., & Wand, M. (2016). Combining Markov random fields and convolutional neural networks for image synthesis. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

[10] Chen, D., Yuan, L., Liao, J., Yu, N., & Hua, G. (2017). Stylebank: An explicit representation for neural image style transfer. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

[11] Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. International Conference on Learning Representations (ICLR).

[12] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR).

[13] Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2018). Progressive growing of GANs for improved quality, stability, and variation. International Conference on Learning Representations (ICLR).

[14] Huang, X., & Belongie, S. (2017). Arbitrary style transfer in real-time with adaptive instance normalization. Proceedings of the IEEE International Conference on Computer Vision.

[15] Kolkin, N., Salavon, J., & Shakhnarovich, G. (2019). Style transfer by relaxed optimal transport and self-similarity. Conference on Computer Vision and Pattern Recognition (CVPR).

[16] Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., & Yang, M. H. (2017). Universal style transfer via feature transforms. Advances in Neural Information Processing Systems (NIPS).

[17] Dumoulin, V., Shlens, J., & Kudlur, M. (2017). A learned representation for artistic style. International Conference on Learning Representations (ICLR).

[18] Selim, A., Elgharib, M., & Doyle, L. (2016). Painting style transfer for head portraits using convolutional neural networks. ACM Transactions on Graphics (TOG), 35(4).

[19] Risser, E., Wilmot, P., & Barnes, C. (2017). Artistic style transfer for videos. Graphical Models, 82, 23-36.

[20] Chen, T. Q., & Schmidt, M. (2016). Fast patch-based style transfer of arbitrary style. arXiv preprint arXiv:1612.04337.

[21] Mechrez, R., Talmi, I., & Zelnik-Manor, L. (2018). The contextual loss for image transformation with non-aligned data. Proceedings of the European Conference on Computer Vision (ECCV).

[22] Sanakoyeu, A., Kotovenko, D., Lang, S., & Ommer, B. (2018). A style-aware content loss for real-time HD style transfer. Proceedings of the EuropeanConference on Computer Vision (ECCV).

[23] Park, D., Liu, S., Wang, T.-C., & Zhu, J.-Y. (2019). Semantic image synthesis with spatially-adaptive normalization. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

[24] Sheng, L., Lin, Z., Shao, J., & Wang, X. (2018). Avatar-net: Multi-scale zero-shot style transfer by feature decoration. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

[25] An, J., & Pellacini, F. (2018). CHAMELEON: Adaptive selection of collections of style-transfer filters. ACM Transactions on Graphics (TOG), 37(4).

[26] Liao, J., Yao, Y., Yuan, L., Hua, G., & Kang, S. B. (2017). Visual attribute transfer through deep image analogy. ACM Transactions on Graphics (TOG), 36(4), 120.

[27] Elgammal, A., Liu, B., Elhoseiny, M., & Mazzone, M. (2017). CAN: Creative adversarial networks, generating "art" by learning about styles and deviating from style norms. International Conference on Computational Creativity.

[28] Gharbi, M., Chen, J., Barron, J. T., Hasinoff, S.., & Durand, F. (2017). Deep bilateral learning for eal-time image enhancement. ACM Transactions on Graphics (TOG), 36(4).

[29] J. Johnson, & Fei-Fei, L. (2018). Image generation from scene graphs. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

[30] Mordvintsev, A., Olah, C., & Tyka, M. (2015). Inceptionism: Going deeper into neural networks. Google Research Blog.