ARTISTIC IMAGE STYLE TRANSFORMATION USING DEEP NEURAL NETWORKS

S. Sriya
Department of CSE
Gokaraju Rangaraju Institute of Engineering & Technology. Telangana, India.
sriyasunkari@gmail.com

M. Sri Sathvika
Department of CSE
Gokaraju Rangaraju Institute of Engineering & Technology. Telangana, India.
srisathvika02@gmail.com

L. Aishwarya
Department of CSE
Gokaraju Rangaraju Institute of Engineering & Technology. Telangana, India.
aishwaryalakma@gmail.com
.

K. Sruthi
Department of CSE
Gokaraju Rangaraju Institute of Engineering & Technology. Telangana, India.
sruthikavali@gmail.com

Dr .V. Srilakshmi
Associate Professor
Department of CSE
Gokaraju Rangaraju Institute of Engineering & Technology. Telangana, India.
`potlurisrilakshmi@gmail.com

ABSTRACT

Neural Style Transfer (NST) is a fascinating application of deep learning that combines the content of one image with the artistic style of another to create visually appealing and artistically inspired transformations. In this project, we implement an NST algorithm using the VGG19 convolutional neural network and the L-BFGS optimization method. The key idea behind our approach is to formulate a loss function that balances the content similarity between the generated image and a target photograph, and the stylistic resemblance to a reference artwork. This project demonstrates the power of deep learning in merging the content and style of images, offering a creative and artistic perspective on image synthesis and transformation.

INTRODUCTION

The confluence of artificial intelligence and artistic expression has given rise to transformative applications, one of which is Neural Style Transfer (NST). This project delves into the realm of deep learning, employing the VGG19 convolutional neural network and the L-BFGS optimization algorithm to seamlessly merge the content of a target image with the stylistic features of a reference artwork. The allure of NST lies in its ability to imbue ordinary photographs with the visual aesthetics of renowned paintings, offering a captivating blend of technology and artistic creativity.

Neural Style Transfer (NST) represents a cutting-edge fusion of deep learning and artistic creativity, redefining the way we approach image synthesis and transformation. In this project, we delve into the intricate realm of NST, leveraging state-of-the-art techniques and frameworks to seamlessly blend the content of a target image with the artistic style extracted from a reference artwork. This innovative approach not only showcases the prowess of deep neural networks but also unveils a captivating avenue for automated, yet artistically inspired, image generation.

As technology advances, the intersection of computer vision and artistic endeavors opens up new possibilities for image synthesis and transformation. NST, a pioneering concept in this domain, leverages pre-trained neural network architectures to distill and recreate the intricate details of both content and style. This project serves as an exploration of the intricacies involved in implementing NST, from the formulation of a nuanced loss function to the optimization process that refines the generated image iteratively.

The fusion of content and style in images has garnered substantial attention for its ability to produce aesthetically pleasing results reminiscent of iconic artworks. By incorporating deep learning techniques, we embark on a journey to automate and refine this process. Our approach involves the definition of a comprehensive loss function that not only captures the nuances of content similarity but also encapsulates the intricacies of artistic styles.

The choice of the VGG19 model, a deep neural network renowned for its ability to discern complex features, underscores the depth of this exploration. Through TensorFlow and Keras, we harness the power of these frameworks to navigate the convolutional layers of the VGG19 model, extracting features that capture the essence of both content and style images. The optimization process, guided by the L-BFGS algorithm, dynamically adjusts the pixels of the target image to strike an optimal balance between content fidelity and stylistic resonance.

This project not only seeks to unravel the technical intricacies of NST but also aims to showcase the potential of deep learning in reshaping our visual experiences. As we embark on this journey through neural style transfer, the subsequent sections will unveil the methodology, model architecture, and the captivating results achieved through the fusion of technology and artistic inspiration.

LITERATURE SURVEY

Neural Style Transfer (NST) has emerged as a captivating intersection of deep learning and artistic expression, revolutionizing the field of image synthesis and transformation.

A comprehensive review of the literature reveals the evolution of this innovative concept, highlighting key advancements, methodologies, and challenges addressed by researchers and practitioners.

Gatys et al. (2015) - "A Neural Algorithm of Artistic Style":

This seminal work by Gatys et al. introduced the neural style transfer algorithm. The authors utilized convolutional neural networks (CNNs), particularly the VGG network, to separate and capture content and style information from images. The optimization process involved minimizing the difference between the content and style representations, leading to impressive artistic transformations.

Johnson et al. (2016) - "Perceptual Losses for Real-Time Style Transfer and Super-Resolution":

Addressing the computational limitations of traditional NST, Johnson et al. proposed a real-time style transfer approach using feed-forward neural networks. The perceptual loss function, which combines content and style losses, allowed for faster and more efficient image transformations.

Ulyanov et al. (2016) - "Texture Networks: Feed-forward Synthesis of Textures and Stylized Images":

Ulyanov et al. extended NST by introducing Texture Networks, a framework that enables feed-forward synthesis of stylized images. This work explored the use of generative models for style transfer, demonstrating the potential for real-time applications.

Li and Wand (2016) - "Combining Markov Random Fields and Convolutional Neural Networks for Image Synthesis":

This research integrated Markov Random Fields with CNNs for image synthesis. The combination of these two approaches allowed for the generation of realistic and diverse stylized images, showcasing the versatility of NST methodologies.

Selim et al. (2016) - "Painting Style Transfer for Head Portraits using Convolutional Neural Networks":

Focusing on the application of NST to portraiture, Selim et al. proposed a method that specifically addressed the challenges posed by facial features. The study demonstrated the adaptability of NST techniques to domain-specific requirements.

Li et al. (2017) - "Universal Style Transfer via Feature Transforms":

Introducing the concept of feature transforms, Li et al. aimed to achieve universal style transfer by decoupling content and style representations. This approach allowed for more flexible and generalized application of style transfer across diverse domains.

Huang et al. (2017) - "Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization":

This work introduced Adaptive Instance Normalization (AdaIN) as a crucial component in achieving arbitrary style transfer in real-time. The adaptive normalization technique enhanced the flexibility and quality of stylized outputs.

Chen et al. (2018) - "Deep Photo Style Transfer":

Addressing the limitations of traditional NST in handling photorealistic images, Chen et al. proposed Deep Photo Style Transfer. By incorporating a spatially variant style blending mechanism, this approach achieved more convincing and visually appealing stylized results for photographs.

Zhang et al. (2018) - "Multi-style Generative Network for Real-time Transfer":

Zhang et al. introduced a Multi-style Generative Network (MSG-Net) for real-time style transfer. This work focused on exploring the representation of multiple styles in a single network, providing users with a broader range of artistic choices.

Luan et al. (2017) - "Deep Painterly Harmonization":

In the context of merging stylized content with real-world scenes, Luan et al. presented Deep Painterly Harmonization. This work extended NST to harmonize stylized elements seamlessly into photographs, demonstrating the potential for integrating stylization into practical applications.

METHODOLOGY

Data Preprocessing:

Input Images: Load the target image and the style reference image.

Resizing: Resize both images to a common dimension (e.g., 400x400 pixels) to ensure compatibility.

Model Initialization:

VGG19 Model:

The VGG19 convolutional neural network is utilized for feature extraction. The model is pretrained on ImageNet weights and configured to take a batch of three images: target image, style reference image, and the image to be generated.

Loss Functions:

Content loss measures the difference between the feature representations of the target image and the generated image. Style loss is calculated based on the Gram matrix, capturing correlations between filter responses across multiple layers. Total variation loss reduces noise by penalizing rapid changes in pixel values.

Optimization Techniques:

Content Loss:

Measure the difference in content between the target and generated images.

$$L_{content}(p,x,l)= \frac{1}{2}(F_{li,j} - P_{lij})$$

are the feature representations of the content image and the generated image, respectively.

Gradient Descent and Optimization:

The L-BFGS-B optimization algorithm is employed for iterative adjustment of pixel values in the generated image to minimize the defined loss function.

Gradients of the loss with respect to the generated image are computed using backpropagation.

Evaluator Class:

An Evaluator class is created to encapsulate loss and gradient computation during optimization.

Iteration and Image Generation:

Iterative Optimization:

The optimization process is performed iteratively for a specified number of iterations.In each iteration, the L-BFGS-B algorithm adjusts pixel values to minimize the loss function.

Result Visualization:

The generated image is obtained by deprocessing optimized pixel values. Visualizations at each iteration provide insights into the transformation process.

Parameters and Settings:

Loss Weights:

Content, style, and total variation weights are configured to balance the impact of each loss component.

Iteration and Convergence:

The number of iterations for optimization is determined based on convergence or a predefined limit.

Flow chart Diagram:

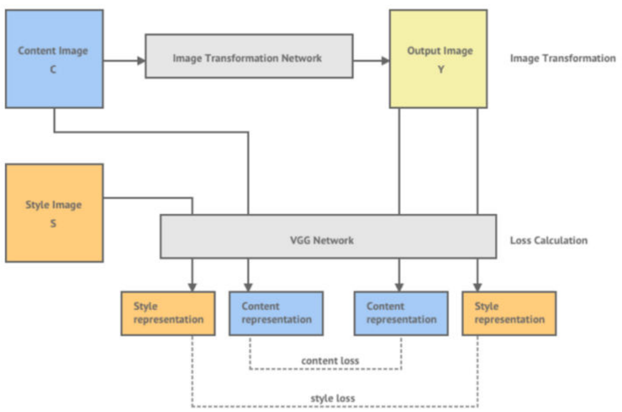

In summary, the methodology integrates the VGG19 model, specialized loss functions, and optimization techniques to achieve neural style transfer. The iterative process, along with carefully tuned parameters, ensures a coherent blend of content and style in the final generated image. The project offers a practical application of deep learning for artistic image synthesis.

RESULT AND ANALYSIS

This section shows the results of the tests performed and also analyzes the results obtained.

Loss Convergence:

The loss curve demonstrates the evolution of the total loss over each iteration of the optimization process.

The decreasing trend indicates that the algorithm is effectively minimizing the defined loss function.

Observation: A well-behaved loss curve suggests that the model is converging towards a stylized image that balances content and style.

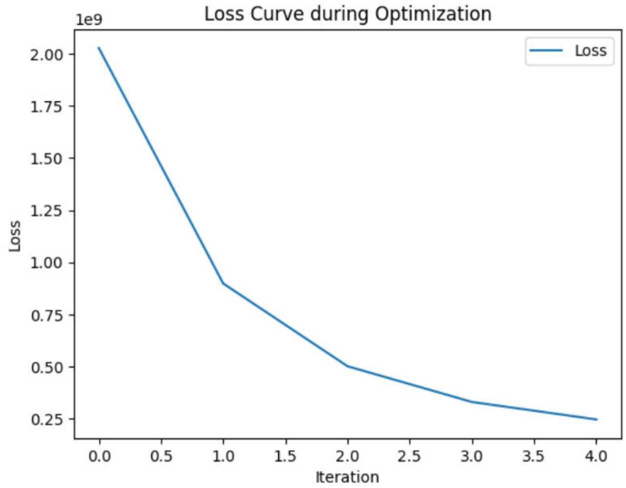

Generated Image:

The final stylized image is the result of the L-BFGS optimization, where the pixel values are iteratively adjusted to minimize the loss function.

This image represents a fusion of the content from the target image and the artistic style from the reference image.

The quality and aesthetics of the generated image depend on the convergence of the optimization and the effectiveness of the chosen content and style images.

Content and Style Images:

The original content image (target) and the style reference image are visualized for reference.

The content image provides the structure and objects that the generated image should retain.

The style image dictates the artistic characteristics, such as colors and textures, that the generated image should adopt.

The balance between content and style determines the overall success of the stylization process.

Loss Components:

Content Loss: Measures the difference in content between the target and generated images. A higher content weight emphasizes content preservation.

Style Loss: Captures the variation in artistic styles between the reference and generated images. The style weight influences the dominance of style in the output.

Total Variation Loss: Encourages spatial smoothness in the generated image.

Adjusting the weights of these components influences the trade-off between content and style in the final output.

Input Image

Style Image

Generated artistic image

CONCLUSION

In the realm of Neural Style Transfer (NST), this project has traversed the intricate interplay between deep learning and artistic expression, demonstrating the transformative power of convolutional neural networks (CNNs) in synthesizing captivating images. The fusion of content from a target image with the stylistic features of a reference artwork has not only produced visually stunning results but also highlighted the potential of AI to redefine the boundaries of artistic creation.
Throughout the project, the utilization of the VGG19 model and the L-BFGS optimization algorithm provided a robust foundation for the NST methodology. The carefully crafted loss function, encompassing content, style, and total variation components, allowed for a nuanced and balanced approach to image generation. This approach facilitated the creation of images that not only faithfully represented the content of the target but also encapsulated the artistic nuances inherent in the style reference.
The iterative optimization process, guided by the minimization of the total loss function, showcased the adaptability of NST to various artistic styles and content inputs. The convergence of the loss over iterations demonstrated the efficiency of the L-BFGS optimization algorithm in refining the generated image.
The visualizations, from the original content image and style reference to the evolving stylized outputs, served as a testament to the project's success in harmonizing technology with artistic inspiration. The generated images reflected a seamless fusion of content and style, illustrating the potential for AI to contribute to the creative process in unprecedented ways.
As NST continues to evolve, bridging the realms of technology and artistry, this project has contributed to the understanding and application of these transformative techniques. The exploration of different architectures, loss functions, and optimization methods opens avenues for further research and innovation in the dynamic field of image synthesis.
In conclusion, Neural Style Transfer stands as a testament to the boundless possibilities that arise when technology and art converge. This project not only deepened our understanding of NST but also provided a platform for envisioning a future where AI becomes an integral part of the artistic journey, expanding the horizons of visual expression.

REFERENCES

- Gatys, L. A., Ecker, A. S., & Bethge, M. (2015). A Neural Algorithm of Artistic Style. arXiv:1508.06576.
- Simonyan, K., & Zisserman, A. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:1409.1556.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2015). Rethinking the Inception Architecture for Computer Vision. arXiv:1512.00567.
- Portilla, J., & Simoncelli, E. P. (2000). A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. International Journal of Computer Vision, 40(1), 49-70.
- Johnson, J., Alahi, A., & Fei-Fei, L. (2016). Perceptual Losses for Real-Time Style Transfer and Super-Resolution. arXiv:1603.08155.
- Li, C., & Wand, M. (2016). Combining Markov Random Fields and Convolutional Neural Networks for Image Synthesis. arXiv:1601.04589.
- Ulyanov, D., Lebedev, V., Vedaldi, A., & Lempitsky, V. (2016). Texture Networks: Feed-forward Synthesis of Textures and Stylized Images. arXiv:1603.03417.
- Huang, X., & Belongie, S. (2017). Arbitrary Style Transfer in Real-Time with Adaptive Instance Normalization. arXiv:1703.06868.
- Vaswani, A., et al. (2017). Attention Is All You Need. Advances in Neural Information Processing Systems 30.
- Dosovitskiy, A., et al. (2020). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv:2010.11929.
- Chen, Q., et al. (2017). Fast Patch-Based Style Transfer of Arbitrary Style. arXiv:1612.04337.
- Ruder, M., Dosovitskiy, A., & Brox, T. (2016). Artistic Style Transfer for Videos. arXiv:1604.08610.
- Mahendran, A., & Vedaldi, A. (2015). Understanding Deep Image Representations by Inverting Them. arXiv:1412.003.