



Enhancing Neural Style Transfer through Integration with Machine Learning Operations Principles

M. Rama Bai¹, D. Deepika², Siddharth Aelpula³, Gandam Revan⁴

Professor, PhD, Department of Emerging Technologies, Mahatma Gandhi Institute of Technology,
Telangana, Hyderabad, India¹

Assistant Professor, MTech, Department of Emerging Technologies, Mahatma Gandhi Institute of Technology,
Telangana, Hyderabad, India²

Student, BTech, Department of Emerging Technologies, Mahatma Gandhi Institute of Technology,
Telangana, Hyderabad, India³

Student, BTech, Department of Emerging Technology, Mahatma Gandhi Institute of Technology,
Telangana, Hyderabad, India⁴

Abstract: This paper presents a comprehensive Machine Learning Operations (MLOps) framework tailored for neural style transfer, focusing on modularity, maintainability, and scalability. Leveraging deep learning models, particularly VGG16, and inspired by seminal works like "A Learned Representation for Artistic Style," our framework integrates cutting-edge MLOps principles to enhance development processes and reproducibility. The implementation utilizes PyTorch for neural networks, FastAPI for backend optimization and MLflow, DVC, and Dagshub for detailed experiment tracking and version control. The frontend is developed with Streamlit, ensuring user-friendly interaction, while Docker guarantees deployment portability. Continuous integration and deployment are managed via GitHub Actions, with AWS ECS and Fargate providing scalability and reliability. Terraform is employed for Infrastructure as Code, enhancing system architecture agility. This end-to-end approach aims to improve model performance, streamline pipelines, and uphold reproducibility and sustainability in neural style transfer applications, pushing the boundaries of innovation in this domain. Our integrated MLOps framework demonstrates significant potential in advancing neural style transfer technology.

Keywords: Machine Learning Operations (MLOps), neural style transfer, deep learning, model deployment, reproducibility, scalability, maintainability

I. INTRODUCTION

Neural Style Transfer (NST) represents a captivating intersection of artificial intelligence and computer vision, aiming to seamlessly merge the content of one image with the artistic style of another. Introduced in the manner of Gatys et al. in their seminal paper "A Neural Algorithm of Artistic Style" in 2016 [9], this innovative approach utilizes convolutional neural networks (CNNs) to disentangle and recombine the content and style attributes of arbitrary images. This paper tries to explore the fusion of Neural Style Transfer with Machine Learning Operations (MLOps) principles, bridging conventional software engineering and machine learning methodologies. MLOps ensures efficient deployment, monitoring, and maintenance of ML models, fostering reproducible development and deployment across domains [2]. It encompasses robust data pipelines, process automation, and continuous integration/continuous deployment (CI/CD) workflows. By integrating NST within the framework of MLOps, we showcase how artificial intelligence techniques can enhance structured operational paradigms, promoting agility, reliability, and scalability in ML model development and deployment. The significance of Neural Style Transfer within the realm of MLOps cannot be overstated. NST's ability to blend the content of one image with the style of another using deep learning algorithms not only demonstrates the power of machine learning but also underscores the importance of efficient operational practices in deploying such advanced techniques in real-world applications. Moreover, MLOps fosters best practices for model management, versioning, and reproducibility, which are critical for ensuring the reliability and maintainability of NST systems over time [3].

With MLOps, organizations can establish robust workflows for model development, testing, and deployment, facilitating seamless collaboration among data scientists, engineers, involved in the Neural Style Transfer pipeline.



Furthermore, recent works such as "Neural Style Transfer for Semantic Domain Translation" by Sanakoyeu et al. ^[1] have highlighted the potential of NST in semantic domain translation, indicating its versatility beyond traditional style transfer tasks. Likewise, Casale et al. ^[2] have proposed "Machine Learning DevOps" as an engineering perspective, shedding light on the operational challenges and solutions in deploying machine learning models at scale. While Neural Style Transfer serves as a compelling showcase for MLOps principles, its integration with efficient operational practices not only enables scalable deployment but also ensures the reliability, fairness, and ethical compliance of AI systems in real-world settings

II. RELATED WORK

Before the implementation of MLOps practices ^[2], the machine learning development and deployment process suffered from fragmentation, manual intervention, and inefficiencies. The lack of a structured approach led to several challenges. Firstly, the absence of a robust versioning system hampered traceability and reproducibility ^[4], making it difficult to track changes in data and models. Additionally, feature extraction methods were suboptimal ^[1], and scaling model training and evaluation was limited due to the absence of efficient distributed computing solutions ^[3]. Moreover, inflexible configuration management hindered adaptability to diverse requirements ^[2] and the integration of advanced models. Model deployment lacked consistency across different environments ^[2], and the existing infrastructure struggled to accommodate increasing workloads efficiently. These shortcomings highlighted the need for a more streamlined and automated approach to machine learning operations. The introduction of MLOps aimed to address these limitations by providing a framework for optimizing the end-to-end ML lifecycle ^[2]. By implementing MLOps, organizations sought to enhance efficiency, consistency, and scalability in model development and deployment processes. While traditional software development follows a structured process akin to baking a familiar cake recipe, MLOps introduces a dynamic and experimental element, akin to adding a secret ingredient that enhances the flavor. It involves integrating data, automation, and continuous learning to ensure the success of machine learning initiatives.

The traditional approach to machine learning development before MLOps was inefficient and fragmented. It lacked structured processes ^[2], leading to manual intervention, inefficiencies, and challenges in version control ^[4], scalability ^[3], and deployment consistency ^[2]. With the adoption of MLOps, organizations transitioned towards a more cohesive and efficient approach to machine learning development and deployment. MLOps practices provided a structured framework for managing the entire ML lifecycle, from data preparation to model deployment and monitoring. Moreover, the integration of MLOps principles facilitated collaboration between data scientists, engineers, and other stakeholders, fostering a more cohesive and agile development environment. By automating repetitive tasks, ensuring version control, and enabling seamless deployment across diverse environments, MLOps practices streamlined the machine learning workflow and improved overall productivity. Overall, the introduction of MLOps marked a significant shift in the machine learning landscape, offering a standardized approach to development and deployment that emphasizes efficiency, scalability, and reproducibility.

III. PROPOSED WORK

This project presents a comprehensive approach to advancing neural style transfer methodologies by integrating MLOps principles into the development pipeline. By leveraging sophisticated deep learning architectures like VGG16, our aim is to foster clarity, maintainability, and sustainability throughout the neural style transfer process. Addressing prevalent challenges in traditional neural style transfer methods, our proposed system emphasizes efficient model integration and streamlined development pipelines. Through the incorporation of MLOps, we envision optimizing the development workflow, enhancing model performance, and ensuring robust reproducibility and scalability. At the core of our system lies PyTorch for neural network implementation, complemented by FastAPI for efficient backend operations. Experiment tracking and version control are meticulously managed using MLflow, DVC, and Dagshub, ensuring transparency and reproducibility of results. For user-friendly visualization, we employ Streamlit, while Docker facilitates seamless deployment and portability across various environments. Automation is key, with GitHub Actions automating testing and deployment processes, ensuring reliability and efficiency. To address scalability requirements, we leverage AWS ECS and Fargate, ensuring the system can handle increasing workloads with ease. Terraform is utilized for managing Infrastructure as Code, enabling agility and ease of management. Through this integrated approach, our project aims to push the boundaries of innovation in neural style transfer applications. By embracing MLOps principles, we seek to elevate model performance, streamline development pipelines, and uphold standards of reproducibility, scalability, and sustainability in the field of neural style transfer.

Through this integrated approach, our project aims to push the boundaries of innovation in neural style transfer applications. By embracing MLOps principles, we seek to elevate model performance, streamline development pipelines, and uphold standards of reproducibility, scalability, and sustainability in the field of neural style transfer.



IV. MLOPS LIFECYCLE IMPLEMENTATION FOR NEURAL STYLE TRANSFER

Machine Learning Operations (MLOps) encompasses practices that govern the full lifecycle of machine learning models, extending beyond development and deployment to include continuous monitoring and updates. As organizations increasingly rely on machine learning for data-driven applications, MLOps has evolved to integrate development activities such as design, build, and testing with operational tasks like deployment, maintenance, and monitoring in a continuous feedback loop. This practice involves collaboration among data scientists, engineers, and IT professionals. Implementing MLOps in neural style transfer projects facilitates the continuous integration and delivery of models, ensuring the application meets user requirements and delivers substantial value. This approach not only optimizes the workflow but also enhances the reliability and scalability of the deployed models.

A. Plan Phase – Strategic Roadmap

For the "Plan" stage of the Neural Style Transfer project, the primary goal is clearly defined, creating a system capable of transferring the artistic style from one image onto the content of another image. The key objectives are to develop either a standalone application or a service that can seamlessly integrate with existing platforms, allowing users to apply various artistic styles to their photos seamlessly. Extensive research is conducted to understand the current landscape, existing solutions, latest advancements in Neural Style Transfer techniques, potential challenges, and best practices. This knowledge informs subsequent decisions and strategies. Data planning involves sourcing relevant datasets for style and content images. A customized version of the MS Coco 2014 train dataset is chosen. Strategies for preprocessing, organizing, and versioning this data are established to ensure consistency and reproducibility during experiments. In terms of technology selection, PyTorch is chosen as the primary machine learning framework due to its flexibility and extensive community support. MLflow is integrated for experiment tracking, and DVC is used for efficient data versioning, enhancing MLOps capabilities. These tools are managed through a single platform called "Dagshub." For the modeling approach, the VGG16 architecture is selected, and appropriate loss functions for content and style are defined, aligning with established best practices in the field. The pipeline design encompasses various stages, including pulling data remotely, model training, evaluation, and deployment. Each step is carefully mapped out to ensure seamless integration and efficient workflow management. A comprehensive testing and monitoring strategy is devised to ensure the system's robustness and reliability. This involves defining key metrics for model evaluation and implementing logging and alerting systems to monitor performance in real-time. Resource allocation is a critical consideration, and the computational resources required for training the model, storing data, and deploying the system are estimated. This allows for optimizing resource utilization and budget effectively. This workflow for this plan is illustrated in Figure 1.




Figure 1 - The Plan Outlook

B. Design Phase – Neural Style Transfer Architecture and Infrastructure

In the design stage of the neural style transfer project, the requirements, architecture, and specifications are carefully crafted to ensure the project's success. The chosen model architecture, VGG16, serves as the foundation for the style transfer network. The requirements stack includes libraries such as Pillow, torch, torchvision, MLFlow, FastAPI, Streamlit, and others, provides the necessary tools for model development, deployment, and monitoring. The architecture of the style transfer network is depicted in Figure 2, illustrating the flow of operations from input to output.

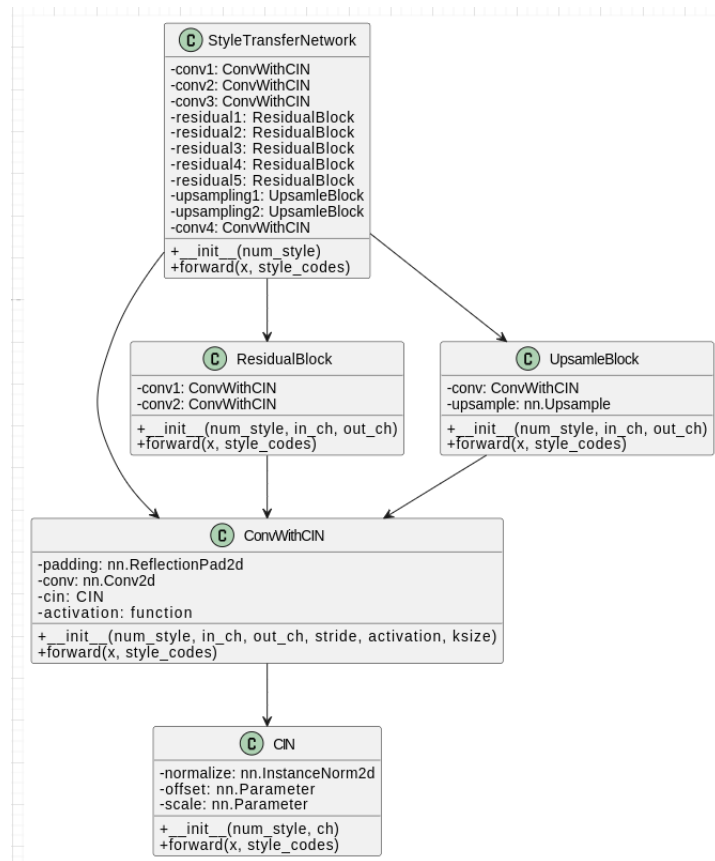


Figure 2 - The architecture of the style transfer network

The neural style transfer architecture consists of convolutional layers, residual blocks, and upsampling blocks to transform content images using style codes, preserving content while applying the style. The data pipeline, defined in the DVC file, manages data pulling, model training, and performance evaluation, ensuring reproducibility and scalability through systematic data versioning and model training. AWS is utilized for hosting, with ECS Fargate and Terraform efficiently managing infrastructure, ensuring scalability and reliability for both backend services and frontend applications. The DVC pipeline is depicted in Figure 3.

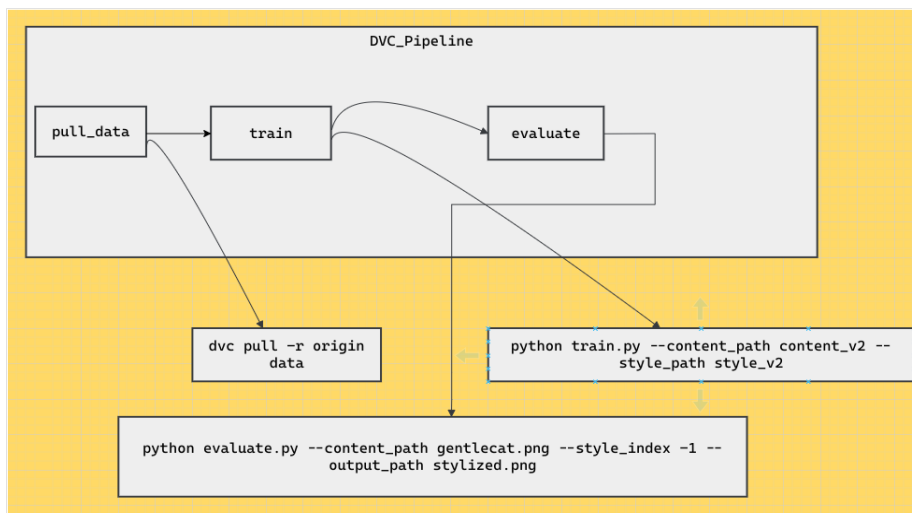


Figure - 3 DVC Pipeline

DVC excels at what git is unable to do, it acts as a wrapper around git by adding the extra functionality of tracking large datasets and thus enabling version control for them as shown in Figure 4

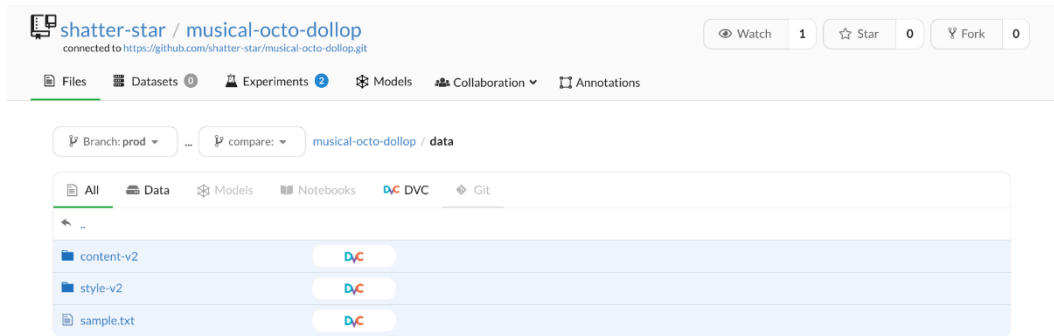


Figure 4 - DVC tracked dataset

C. Build Phase – Optimizing Neural Style Transfer Build Process

The training process is meticulously defined to ensure the refinement and optimization of the Neural Style Transfer model. Within this process, several essential steps are undertaken. Initially, the training procedure involves setting up MLflow tracking, a critical component for monitoring and logging experiment metrics. This encompasses defining the MLflow tracking URI and credentials, ensuring seamless integration with the MLflow environment. Following this, data loading and processing are executed. This phase is facilitated by dedicated classes responsible for creating dataloaders for both content and style images, thereby streamlining data ingestion and preprocessing tasks. Furthermore, the training script handles environment setup responsibilities, ensuring the requisite directories are accessible for smooth module access. Parameter logging emerges as another crucial facet, with the training script meticulously recording essential parameters such as style weight, total variation weight, learning rate, batch size, and iterations. This comprehensive logging mechanism not only facilitates experiment reproducibility but also provides invaluable insights into model performance and behaviour over time. In addition to parameter logging, the training script optimizes GPU utilization, leveraging the computational power of multiple GPUs if available. Through efficient parallel processing, training efficiency and throughput are significantly enhanced. Finally, experiment tracking remains paramount, with MLflow serving as the central hub for logging experiment parameters and losses. By systematically recording these metrics, the training script enables comprehensive analysis and evaluation of model performance, driving iterative improvements and advancements in the Neural Style Transfer model. The process of the build stage and experiment tracking, model versioning is shown in Figure 5 and 6.




Figure 5 - Flow of Build Stage

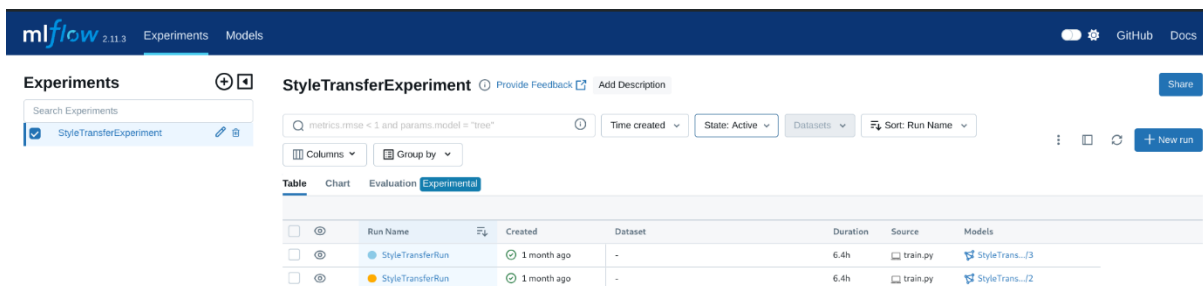


Figure 6 - Experiment Tracking with MLFLOW



The Parameters which are logged are shown in the Table 1

Parameters (5)

Parameter	Value
style_weight	5.0
tv_weight	1e-05
learning_rate	0.0001
batch_size	8
iterations	25000

Table 1 - Parameters logging

D. Test Phase – Ensuring Reliability in Model Components

This suite of tests plays a crucial role in verifying the functionality and correctness of essential components within the Neural Style Transfer project. Covering a range of functionalities including image utilities, data handling, and loss calculations, these tests ensure the robustness and reliability of the codebase. For instance, they verify the accuracy of image loading, saving, and transformations, confirming that images are correctly processed and saved in various formats. Additionally, they validate the dataset class, ensuring that it accurately reports the number of images and returns the correct images and indices when indexed. Furthermore, the tests assess the performance of loss functions, confirming that content, style, and total variation losses are calculated accurately. By running these tests regularly, developers can detect and address any bugs or regressions early on, thereby enhancing the overall quality and maintainability of the codebase.

E. CI/CD Phase – Automation with GitHub Actions and Terraform

CI/CD, which stands for continuous integration and continuous delivery/deployment, aims to streamline and accelerate the software development lifecycle. In the project this is mainly carried out with the help of GitHub actions which automates the process and is complimented using Terraform to provision resources in AWS for a more productive workflow. The cycle is illustrated in Figure 7.

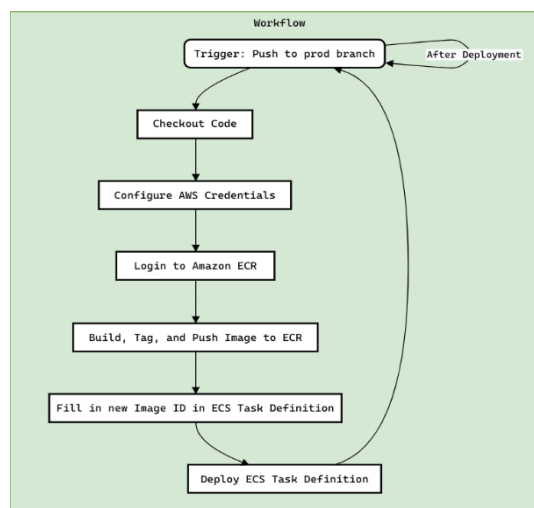


Figure 7 - CI/CD Pipeline

When push events occur in the 'prod' branch, the deployment pipeline starts, setting environment variables for AWS and ECS. The 'Deploy' job runs on an Ubuntu platform, checking out code, configuring AWS credentials with GitHub secrets, and pushing Docker images to Amazon ECR. The image URI is updated in the ECS task definition and deployed to the ECS service, followed by a stability check. Each successful deployment can trigger the next pipeline iteration. Terraform provisions a VPC with public subnets, configures security groups and internet access, and creates an ECR repository



"app_repo" for Docker images. An ECS cluster "app_cluster" and service "app_service" run on EC2 instances, pulling images from "app_repo" and listening on port 80. This setup provides a robust AWS infrastructure for containerized applications, as shown in Figures 8.

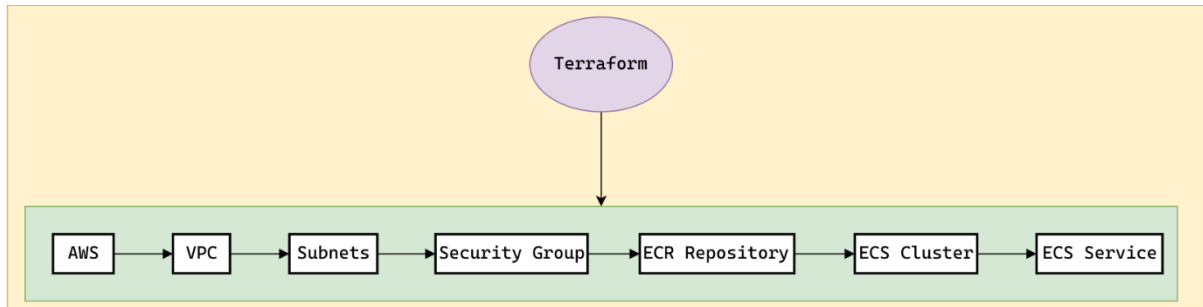


Figure 8 - Schema for Deploying resources on AWS using Terraform

F. Deployment Phase – Launch with FastAPI and AWS ECS

At the core of the project is a FastAPI-based style transfer service that provides an API endpoint /stylize. This endpoint accepts a content image and a style index, using a pre-trained StyleTransferNetwork model from an MLflow model registry for stylization. The application is containerized with Docker for consistent deployment across environments. The deployment is automated using GitHub Actions, and Amazon ECS handles the container orchestration. The ECS task definition, detailed in td.json, specifies the Docker image, CPU, memory, and network settings. Utilizing AWS ECS Fargate and ECR, the project achieves scalability and cost-effectiveness. This deployment strategy, combining FastAPI, Docker, and AWS, ensures efficient, reliable, and up-to-date machine learning applications. The API specs for deployment via Dockerfile are outlined in Table 2.

Table 2 - Deployment ready API Specifications

Name of API	Request Type	URL	Data(sample)	Response(sample)	Description
stylize	POST	/stylize	content_image: File style_index: int	Streaming response with the stylized image	Applies style transfer to a content image using a pre-trained model, uploads the result to S3, and returns it as a streaming response.
root	GET	/	-	{"message": "Style Transfer API is running!"}	Root endpoint that returns a simple message indicating that the API is running.

The AWS deployment is configured for the ap-south-2 region (Asia Pacific - Mumbai) and sets up a VPC with a 10.0.0.0/16 CIDR block and three public subnets across different availability zones. A security group allows inbound traffic on ports 80 and 443 and all outbound traffic. An internet gateway enables internet access, with appropriately configured route tables. An ECR repository named "app_repo" stores Docker images, with image scanning disabled for efficiency. An ECS cluster named "app_cluster" and a service named "app_service" are deployed on EC2 instances. The service's task definition specifies container resources, including CPU, memory, and port mappings, with defined task and execution roles for seamless AWS integration. The deployed container in ECS is shown in Figure 9, and the API can be verified via the ECS service's public IP address (Figure 10).

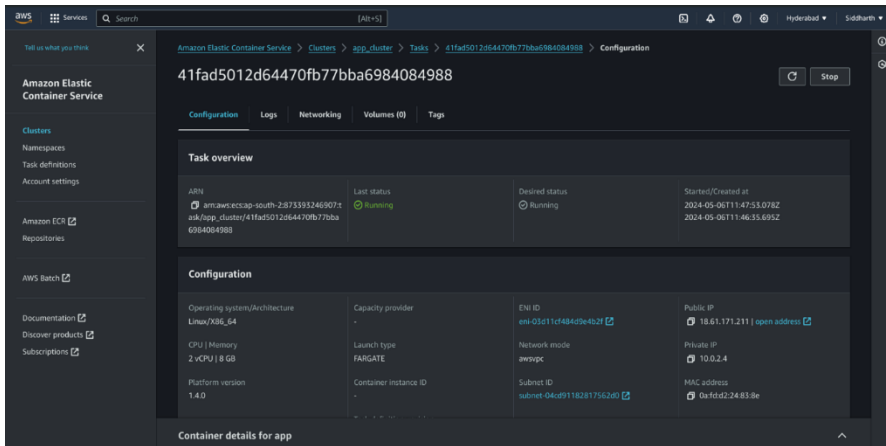


Figure 9 - Deployed ECS cluster on AWS

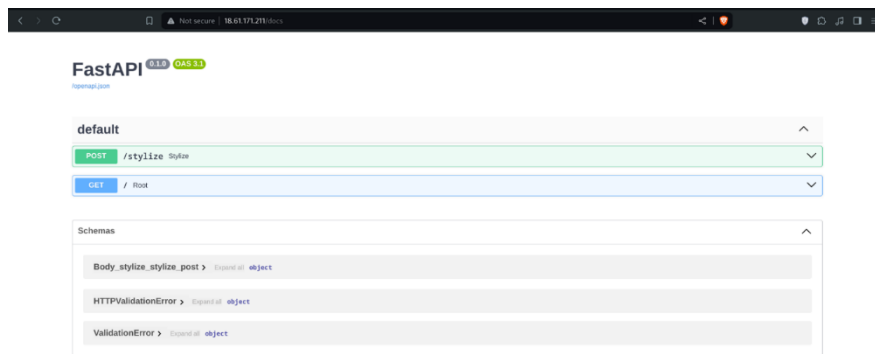


Figure 10 - Successful Deployment of FastApi to AWS ECS

The deployment and CI/CD pipeline primarily support the FastAPI backend. Once deployed, the public IP address can be used to access the /stylize endpoint, integrated into the Streamlit-based frontend. This setup technically constitutes a multi-cloud deployment, with the frontend hosted on Streamlit Community Cloud. The application UI is shown in Figure 11.

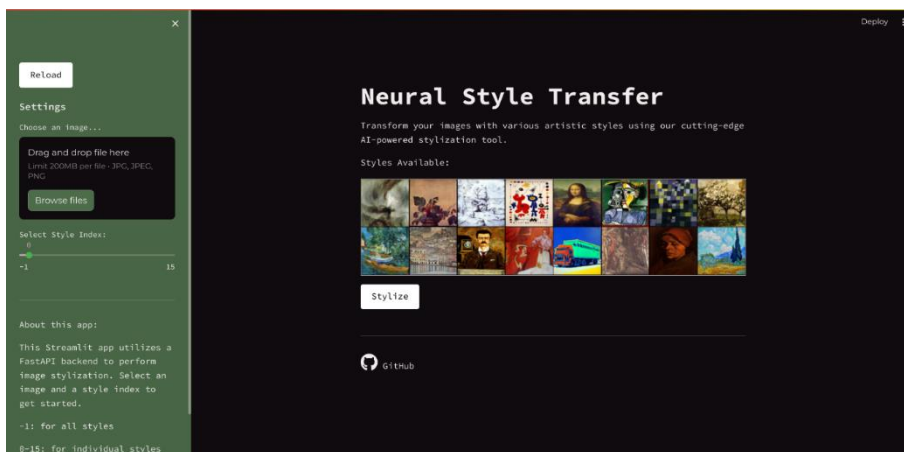


Figure 11 - Streamlit Application

G. Monitor and Operate Phase – Observing model and System efficacy

To ensure effective monitoring for a neural style transfer application, various tools and strategies are used. Resource Utilization Monitoring with apitaly tracks CPU, GPU, and memory usage to optimize infrastructure and costs. Model performance and drift monitoring are managed by MLflow, which continuously evaluates metrics like loss



function values, complemented by human evaluation. Versioning and Logging through MLflow ensure reproducibility and compliance. Model Deployment and Scaling are facilitated by MLflow, DVC, and AWS ECS Fargate, allowing seamless deployment and automated scaling for maintaining performance and reliability. Observability details are shown in Figure 12. Efficient operations on Amazon ECS involve containerization, task definition, cluster configuration, service setup, monitoring, security, CI/CD automation, and cost optimization. This systematic approach ensures streamlined deployment, scalability, performance, and resource management. Prioritizing operational excellence allows for reliable, flexible, and cost-effective execution of neural style transfer models in the cloud. If model or data drift is detected, the model can be seamlessly updated in the production environment after thorough analysis, leading to reassessment and process iteration.

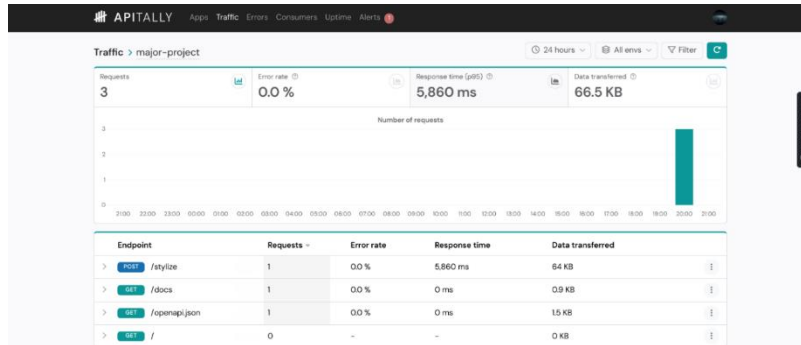


Figure 12 - Administering the API & Monitoring Model Drift

V. TESTING AND RESULTS

In the "Testing Approach" for the neural style transfer model, the evaluation primarily focuses on analysing the performance of the model through the examination of key loss components – content loss, style loss, TV loss, and total loss – across multiple training iterations. The approach involves, Initial Loss Analysis, the testing commences with an assessment of the initial loss values for each component. This stage serves to understand the starting point of the model's training process and provides insights into the challenges associated with content preservation and style transfer. Temporal Analysis, throughout the training iterations, continuous monitoring of the loss curves allows for a temporal analysis of the model's learning process. Key observations include the rapid decrease in content and style losses during the initial training phase, indicating the model's capability to efficiently capture content features and artistic styles from input images. Stability Evaluation, as the training progresses, attention is given to the stability and convergence of the loss curves. Notably, the stabilization of content and style losses at relatively low values signifies the model's ability to maintain content fidelity and transfer styles consistently over time. Investigating the impact of total variation regularization on image smoothness, fluctuations in the TV loss curve prompt an analysis. Despite these fluctuations, manageable TV loss magnitudes indicate the model's ability to generate coherent stylized outputs. Furthermore, the convergence of all loss components at the end of training suggests the model's stability and effective learning. This comprehensive testing approach provides valuable insights into the model's performance in content preservation and style transfer tasks, with key metrics summarized in Table 3 post-training.

Metrics (4)

Metric	Value
content_loss	18.98065948486328
style_loss	2.525583267211914
total_loss	31.60858154296875
tv_loss	0.6502508521080017

Table 3 – Metric Table

The content loss starts extremely high but rapidly decreases within the first few thousand iterations. After that initial drop, it remains relatively low and stable, indicating that the model can preserve the content of the input images reasonably well after the early stages of training, the style loss follows a similar pattern to the content loss, starting high but decreasing rapidly in the initial iterations. However, it stabilizes at a slightly higher value compared to the content



loss, suggesting that transferring the artistic style is a more challenging aspect of the task. Also, the TV loss, which acts as a regularizer for spatial smoothness, starts relatively low but exhibits some fluctuations throughout the training process. This could indicate that the model occasionally struggles to produce spatially coherent stylized images, but the overall magnitude of the TV loss remains manageable. Finally, the total loss, being the combination of the other losses, follows a similar trend to the content and style losses. It starts extremely high but decreases rapidly in the early stages of training, eventually stabilizing at a moderate value. The convergence Plot is shown in Figure 13.

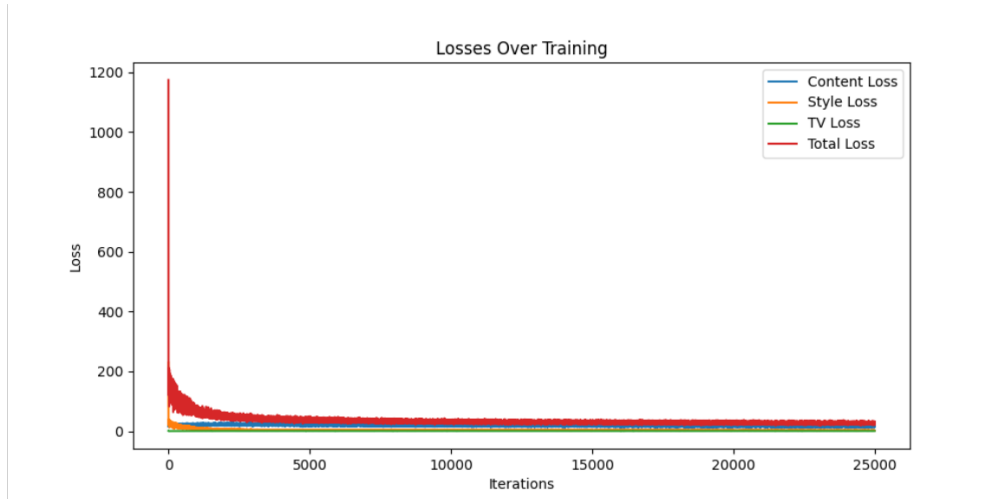


Figure 13 - Plot for convergence

Convergence is evident as all loss curves stabilize towards the end of training, indicating the model has likely mastered the task given the current setup, as shown in Figure 13. The result, a Streamlit web application enables users to upload images and select a style index. Upon clicking "Stylize", the app sends a POST request to the FastAPI backend, which handles the following operations: saving the uploaded content image, loading the pre-trained StyleTransferNetwork model from an MLflow repository, applying the style transfer algorithm, and returning the stylized image as a streaming response. This action is shown in Figure 14.

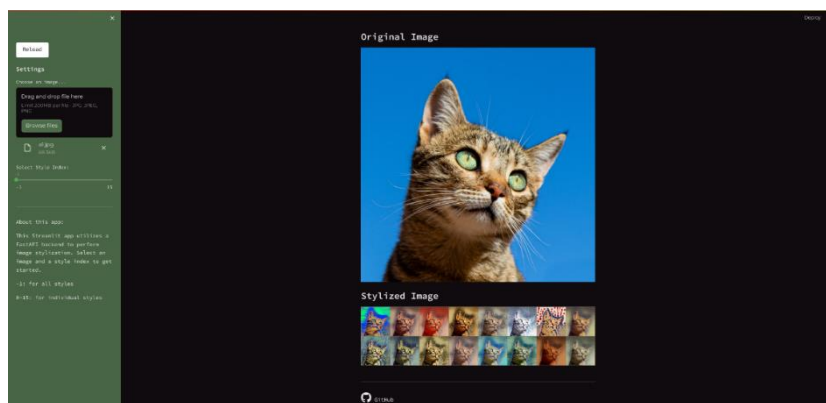


Figure 14 – Result Case: For all Styles

The Streamlit app displays the stylized image alongside the original uploaded image. Key components include the Streamlit web app for the interface, the FastAPI backend for style transfer logic, and a pre-trained StyleTransferNetwork model from an MLflow repository. The integration with an AWS S3 bucket for storing stylized images and a streaming response is generated via the S3 URL. There are a total of 16 style presets, the user can choose any of the corresponding index of the desired style or select all styles at once using the specified style codes as shown in Figure 15 respectively.

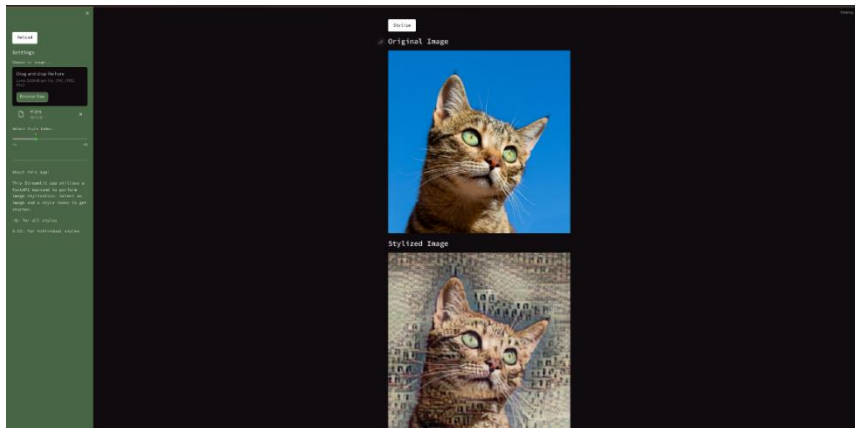


Figure 15 – Result Case: For Individual Style

VI. CONCLUSION

In conclusion, this neural style transfer project stands as a testament to the transformative potential of MLOps principles in modern machine learning workflows. By leveraging automation, collaboration, and scalability, it has redefined the paradigm of image stylization. Through the seamless integration of CI/CD pipelines and a forward-thinking multi-cloud strategy, this project exemplifies the efficiency and agility of contemporary software development practices. The final losses (content: 18.98, style: 2.53, total variation: 0.66, total: 31.61) demonstrate the effective ability and convergence of the model, indicating a promising sign of its efficacy in production. The adoption of best practices in MLOps standards ensures meticulous tracking and management of model iterations, fostering transparency and reproducibility in the deployment and development processes.

Looking ahead, the impact of this project transcends its immediate application, paving the way for significant advancements in machine learning workflows. By demonstrating the effectiveness of MLOps principles in driving innovation and efficiency, it serves as a blueprint for future developments across diverse domains. From enhancing image processing techniques to optimizing complex algorithms, the lessons learned from this project will inform and inspire the next generation of AI solutions. As organizations continue to embrace the potential of MLOps, the ripple effects of this project will be felt far and wide, reshaping the landscape of machine learning and propelling us towards a future where intelligent systems play an increasingly integral role in our lives.

REFERENCES

- [1] Sanakoyeu, A., Tjandraatmadja, C., Kokkinos, I., and Vedaldi, A. (2022). Neural Style Transfer for Semantic Domain Translation. *IEEE Transactions on Pattern Analysis and MI*.
- [2] Casale, G., Munshi, N., and Memoli, P. (2022). Machine Learning DevOps: An Engineering Perspective. *IEEE Transactions on Software Engineering*.
- [3] Jiang, Y., Perkins, H., Liao, S., Zhang, Y., Zhu, J., and Tao, C. (2021). AutoMLOps: Automated Machine Learning Operations. *IEEE International Conference on Data Engineering*
- [4] Paleyes, A., Urma, R.G., and Holden, N. (2020). The Challenges of Continual Learning for Production Machine Learning Pipelines. *ICML Workshop on Continual Learning*.
- [5] Neural Style Transfer: A review (Jing Y. et al., 2019)
- [6] Sato, D., Bhole, A., Cannane, E., and Erguler, D. (2019). Towards DevOps for Machine Learning. *IEEE International Conference on Cloud Computing (CLOUD)*.
- [7] Huang, X. and Belongie, S. (2017). Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization. *IEEE International Conference on Computer Vision (ICCV)*.
- [8] Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., and Yang, M.H. (2017). Universal Style Transfer via Feature Transforms. *Advances in Neural Information Processing Systems (NeurIPS)*.
- [9] Gatys, L.A., Ecker, A.S., and Bethge, M. (2016). Image Style Transfer Using Convolutional Neural Networks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [10] A Learned Representation for Artistic Style (Dumoulin et al., 2016)