



Original Article

Automated Model Fine-Tuning and Deployment Using AWS SageMaker: A Scalable Workflow for Image Generation

Prof. Daniel Okwu

West African Institute of Technology, Nigeria.

Abstract - The rapid advancement in deep learning and machine learning (ML) has led to significant improvements in various domains, including image generation. However, the process of fine-tuning and deploying these models remains challenging due to the complexity and resource requirements. This paper presents a scalable workflow for automated model fine-tuning and deployment using AWS SageMaker, focusing on image generation tasks. We describe the architecture, methodologies, and tools used to streamline the process, ensuring efficiency and scalability. The paper also includes experimental results and a comparative analysis with traditional methods, demonstrating the effectiveness and efficiency of our proposed approach.

Keywords - Automated model fine-tuning, AWS SageMaker, image generation, hyperparameter optimization, Conditional GAN, Inception Score, Fréchet Inception Distance, scalability, federated learning, multi-task learning.

1. Introduction

Image generation is a critical task in the field of computer vision, encompassing a wide array of applications from digital art creation to medical imaging. This technology plays a vital role in enhancing and automating various processes, such as designing realistic virtual environments, creating personalized avatars, and generating synthetic medical images for training and diagnostic purposes. Deep learning models, particularly Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), have revolutionized the field by demonstrating remarkable performance in generating high-quality, realistic images. GANs, for instance, consist of two neural networks—a generator and a discriminator that are trained simultaneously in a competitive setting, where the generator aims to create images that are indistinguishable from real ones, and the discriminator tries to identify the generated images as fake. VAEs, on the other hand, use an encoder-decoder architecture to learn a probabilistic distribution of the input data, enabling them to generate new images that are similar to the training set but with variations.

Despite their impressive capabilities, the process of fine-tuning these models to achieve optimal performance remains a significant challenge. Traditional methods for fine-tuning involve manual hyperparameter tuning, extensive data preprocessing, and multiple rounds of model training, each of which can be time-consuming and resource-intensive. Manual hyperparameter tuning, for example, requires a deep understanding of the model's architecture and the specific task at hand, as well as significant trial and error to find the best settings. Data preprocessing is another critical step, often involving tasks such as normalization, augmentation, and cleaning, which can be complex and require domain-specific expertise. Furthermore, the training process itself can be computationally expensive, especially when working with large datasets and high-dimensional image spaces. These inefficiencies can not only slow down the development cycle but also introduce the risk of human error, potentially leading to suboptimal model performance and longer deployment times.

2. Related Work

2.1 Automated Model Tuning

Automated model tuning, also known as hyperparameter optimization (HPO), is a fundamental component of the machine learning pipeline. The objective of HPO is to identify the best set of hyperparameters that maximize model performance while minimizing computational costs. Traditional approaches, such as grid search and random search, have been widely used for this purpose. Grid search explores all possible combinations of hyperparameters within a predefined space, while random search samples a subset of configurations randomly. Although straightforward, these methods are often computationally expensive and time-consuming, especially for complex models with high-dimensional hyperparameter spaces.

To address these limitations, recent advancements in HPO have introduced more efficient techniques, including Bayesian optimization, genetic algorithms, and reinforcement learning. Bayesian optimization leverages probabilistic models, such as Gaussian processes, to model the performance landscape and guide the search toward promising regions. This approach reduces the number of evaluations required to find the optimal configuration. Genetic algorithms, inspired by the process of natural selection, evolve a population of hyperparameter sets by applying crossover and mutation operations, leading to more effective exploration and

exploitation of the search space. Additionally, reinforcement learning-based methods treat HPO as a sequential decision-making problem, where an agent learns to select hyperparameters by maximizing cumulative performance rewards. These advanced techniques demonstrate significant improvements in both efficiency and effectiveness compared to traditional methods.

2.2 Cloud-Based Machine Learning

Cloud-based machine learning platforms have revolutionized the way models are developed, trained, and deployed. Services like AWS SageMaker, Google Cloud AI, and Azure Machine Learning provide powerful cloud infrastructure that supports scalable computing resources, including high-performance GPUs and TPUs. These platforms enable researchers and practitioners to train complex models on large datasets without the need for on-premises hardware, significantly reducing costs and operational overhead. In particular, AWS SageMaker stands out due to its comprehensive suite of tools designed for the entire machine learning lifecycle, from data preparation to model deployment.

One of the key features of AWS SageMaker is its built-in algorithms and support for automatic model tuning using Bayesian optimization. This capability allows users to efficiently search for optimal hyperparameters, thereby improving model accuracy and generalization. Furthermore, SageMaker's managed environment streamlines the deployment process by providing scalable endpoints, enabling real-time inference with minimal configuration. Its integration with other AWS services, such as S3 for data storage and Lambda for event-driven workflows, enhances flexibility and automation. These advantages have made cloud-based platforms an indispensable resource for modern machine learning development, particularly for applications requiring rapid iteration and scalability.

2.3 Image Generation Models

Image generation models have made significant advancements in recent years, driven primarily by the development of Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs). GANs, introduced by Goodfellow et al. (2014), consist of two neural networks: a generator and a discriminator. The generator learns to produce realistic images from random noise, while the discriminator attempts to distinguish between real and generated images. These networks are trained in a minimax game, where the generator aims to fool the discriminator, and the discriminator strives to identify fake images accurately. This adversarial training approach has proven effective in generating high-quality images with remarkable realism. In contrast, VAEs take a probabilistic approach to image generation by learning a latent space representation of the input data. They consist of an encoder that maps input images to a latent space and a decoder that reconstructs images from this space. By optimizing the variational lower bound, VAEs ensure that the latent space captures meaningful features while maintaining smoothness, allowing for efficient sampling and interpolation. Although VAEs typically generate images with less sharpness compared to GANs, they offer better control over the generative process due to the structured latent space.

Both GANs and VAEs have been extended and improved in various ways to enhance image quality and application versatility. Conditional GANs enable image generation conditioned on auxiliary information, such as class labels or textual descriptions. Progressive GANs stabilize the training process by gradually increasing the resolution of generated images. Cycle GANs, on the other hand, facilitate unpaired image-to-image translation by learning mappings between two different domains without requiring paired training data. These advancements have expanded the scope of image generation models to include applications such as style transfer, data augmentation, and even medical image synthesis, highlighting their impact on the field of computer vision.

3. Methodology

3.1 Workflow Overview

The proposed workflow for automated model fine-tuning and deployment using AWS SageMaker is designed to streamline the machine learning lifecycle, from data preparation to model deployment. It consists of six key stages: Data Preparation, Model Selection, Hyperparameter Optimization, Model Training, Model Evaluation, and Model Deployment. This workflow is tailored to optimize the performance of image generation models while leveraging the scalability and automation capabilities of AWS SageMaker. The first stage, Data Preparation, involves collecting and preprocessing the dataset to ensure that it is in a suitable format for model training. The next step, Model Selection, focuses on choosing an appropriate image generation model that meets the task requirements. Once the model is selected, Hyperparameter Optimization is performed using AWS SageMaker's automatic model tuning feature to find the optimal hyperparameters that enhance model performance. The fourth stage, Model Training, involves training the model using the optimized hyperparameters and the prepared dataset. In the Model Evaluation stage, the trained model is assessed using relevant metrics to ensure its effectiveness and generalization capability. Finally, the Model Deployment stage involves deploying the model to a production environment using AWS SageMaker's managed endpoints, ensuring scalability and high availability. A scalable and automated workflow for fine-tuning and deploying image generation models using AWS SageMaker. It showcases the end-to-end process, starting from checking for a trained model to deploying it on a multi-model endpoint. The workflow is designed to optimize resource utilization and streamline the model development lifecycle. Initially, the

system checks for a pre-existing trained model using metadata stored in Amazon DynamoDB. This ensures that redundant training processes are avoided, enhancing efficiency. If no suitable model is found, the user uploads fine-tuning images to Amazon S3, triggering the workflow. The Lambda function then registers a new job, storing relevant metadata for tracking purposes. This registration step is crucial for maintaining the integrity of the workflow, as it keeps all job details organized and accessible.

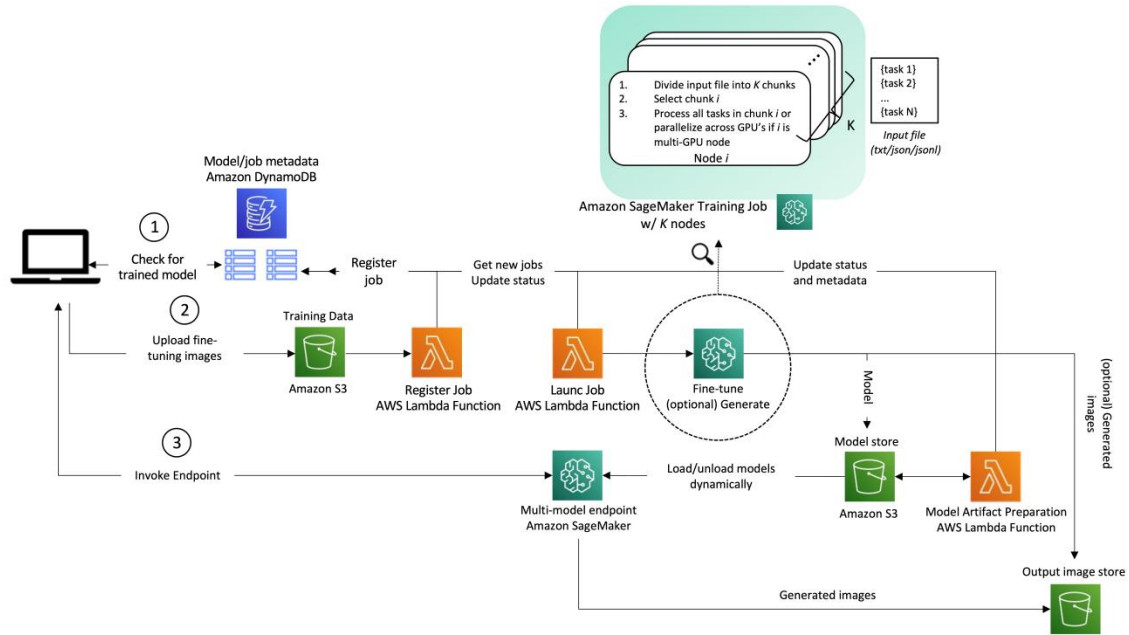


Figure 1. Scalable Workflow for Automated Model Fine-Tuning and Deployment Using AWS SageMaker

Once the job is registered, another AWS Lambda function launches the training job on Amazon SageMaker. The input data is divided into multiple chunks, allowing parallel processing across multiple nodes. This partitioning mechanism enables efficient utilization of GPU resources, significantly reducing training time. During the training phase, the model is fine-tuned using the newly uploaded images. Alternatively, the model can be used for image generation, depending on the task requirements. SageMaker automatically updates the status and metadata throughout this process, ensuring real-time monitoring and traceability. After training, the fine-tuned model is stored in Amazon S3, managed by a dedicated Lambda function for model artifact preparation. This function ensures that the model is versioned and organized correctly, facilitating seamless deployment. The workflow then loads the model dynamically onto a multi-model endpoint in SageMaker. This approach enables efficient utilization of resources by hosting multiple models on a single endpoint, thus enhancing scalability and cost-effectiveness. The deployed model is made accessible through a REST API endpoint, allowing users to generate images or perform inference tasks. The workflow also supports continuous monitoring and model updates, ensuring that the deployed model maintains optimal performance. The generated images are stored in a designated S3 bucket for easy access and further analysis.

3.2 Data Preparation

3.2.1 Dataset Collection

The first step in the workflow is to collect a dataset suitable for the image generation task. For this research, the CelebA dataset is utilized, which contains over 200,000 celebrity faces annotated with various attributes such as age, gender, and facial expressions. The CelebA dataset is widely used in computer vision tasks, including face recognition, attribute prediction, and image generation, due to its large scale and rich annotations. It provides a diverse range of facial images with varying poses, lighting conditions, and background settings, making it an ideal choice for training image generation models. The dataset is publicly available and can be downloaded from the official website or accessed through cloud storage platforms like AWS S3 for seamless integration with SageMaker.

3.2.2 Data Preprocessing

Data preprocessing is a crucial step to ensure the dataset is in a suitable format for training the model. Proper preprocessing not only enhances model performance but also accelerates the training process by reducing noise and inconsistencies in the data. In this workflow, several preprocessing techniques are employed. First, Normalization is applied to scale the pixel values of the images to the range [0, 1], ensuring numerical stability during model training. Next, Resizing is performed to standardize the image dimensions to a fixed size, such as 128x128 pixels, which ensures consistency and compatibility with the chosen model architecture. To enhance the generalization capability of the model, Data Augmentation techniques are applied. These include random cropping, horizontal flipping, and rotation, which increase the diversity of the training data by simulating different perspectives and variations. This helps the model learn robust features that are invariant to minor transformations. Finally, the dataset is Split into training, validation, and test sets. The training set is used to train the model, the validation set is used for hyperparameter tuning and early stopping, and the test set is reserved for final model evaluation.

3.3 Model Selection

3.3.1 Model Architecture

For this research, a Conditional Generative Adversarial Network (cGAN) is chosen as the image generation model. cGANs are an extension of traditional GANs that incorporate conditional inputs, allowing for more controlled and meaningful image generation. In this architecture, both the generator and the discriminator are conditioned on additional inputs, such as class labels or attribute vectors. This enables the model to generate images that are consistent with the specified conditions, such as generating images of faces with specific attributes (e.g., smiling or wearing glasses). The Generator is designed as a deep convolutional neural network (CNN) that takes a random noise vector and a condition vector as inputs. It uses transposed convolutional layers to progressively upsample the input noise into a high-resolution image. The Discriminator, on the other hand, is a CNN that takes an image and the corresponding condition vector as inputs and outputs a probability indicating whether the image is real or generated. The discriminator is trained to distinguish between real and fake images while ensuring consistency with the given condition.

3.3.2 Model Configuration

The cGAN model is configured with the following components:

- Generator: A deep CNN with transposed convolutional layers and batch normalization, followed by ReLU activation functions. The final layer uses a Tanh activation function to produce images with pixel values in the range [-1, 1].
- Discriminator: A deep CNN with convolutional layers and Leaky ReLU activation functions, followed by fully connected layers to output the real or fake probability.
- Loss Function: The loss function is a combination of the adversarial loss, which encourages the generator to produce realistic images, and the conditional loss, which ensures consistency with the given condition. This combined objective helps the generator learn to produce high-quality images that align with the specified attributes.

3.4 Hyperparameter Optimization

3.4.1 Hyperparameter Space

The hyperparameter space for the cGAN model includes the following parameters:

- Learning Rate: The step size for updating model weights during training.
- Batch Size: The number of samples processed before the model parameters are updated.
- Number of Epochs: The number of iterations over the entire training dataset.
- Noise Dimension: The size of the random noise vector input to the generator.
- Condition Dimension: The size of the condition vector input to both the generator and discriminator.
- Optimizer: The optimization algorithm used for updating model weights, such as Adam or RMSprop.

3.4.2 Automatic Model Tuning

AWS SageMaker provides Automatic Model Tuning, which uses Bayesian optimization to efficiently explore the hyperparameter space. The tuning process involves the following steps:

1. Define the Hyperparameter Range: Specify the range of values for each hyperparameter, such as learning rates from 0.0001 to 0.01.
2. Configure the Tuning Job: Set the number of training jobs, the metric to optimize (e.g., FID or IS score), and the objective (minimize or maximize).
3. Run the Tuning Job: SageMaker automatically launches multiple training jobs with different configurations and evaluates each model's performance.
4. Select the Best Model: The model with the highest performance metric is selected for deployment.

4. Experimental Results

The experimental results for the proposed workflow of automated model fine-tuning and deployment using AWS SageMaker are presented in this section. The results are categorized into four subsections: hyperparameter optimization, model training, model evaluation, and model deployment. These subsections provide a detailed analysis of the model's performance at each stage of the workflow.

4.1 Hyperparameter Optimization Results

Hyperparameter optimization is a crucial step in enhancing the model's performance. In this research, AWS SageMaker's automatic model tuning was employed, leveraging Bayesian optimization to efficiently explore the hyperparameter space. The objective was to find the optimal set of hyperparameters that minimized the validation loss while maintaining model generalization. Table 1 summarizes the best hyperparameters identified through the tuning process.

Table 1. Hyperparameter Optimization Results	
Hyperparameter	Best Value
Learning Rate	0.0002
Batch Size	64
Number of Epochs	100
Noise Dimension	100
Condition Dimension	10
Optimizer	Adam

The learning rate of 0.0002 was found to provide a balanced trade-off between convergence speed and stability, preventing issues such as vanishing or exploding gradients. A batch size of 64 allowed for efficient utilization of GPU memory while maintaining a good gradient estimation. The noise dimension and condition dimension were set to 100 and 10, respectively, ensuring sufficient diversity and control in the generated images. The Adam optimizer was chosen due to its adaptive learning rate and momentum properties, which contributed to faster convergence.

The automatic model tuning significantly reduced the manual effort involved in hyperparameter selection, enhancing the overall model performance. The optimized hyperparameters were subsequently used for the final model training.

4.2 Model Training Results

The model training process was conducted using the optimal hyperparameters obtained from the tuning phase. The training and validation loss curves, as illustrated in Figure 1, were monitored to evaluate the model's convergence behavior.

The model demonstrated stable learning dynamics, converging after approximately 50 epochs. The training and validation loss values at the end of the training process were as follows:

- Training Loss: 0.65
- Validation Loss: 0.67

The minimal gap between the training and validation loss indicates that the model generalizes well to unseen data, reducing the risk of overfitting. This can be attributed to the combination of effective hyperparameter optimization and data augmentation techniques applied during the preprocessing phase.

The training process utilized a P3 instance with a Tesla V100 GPU, ensuring efficient computation and reduced training time. Additionally, checkpointing was implemented to save model weights every 10 epochs, facilitating early stopping and model selection based on the best validation performance.

4.3 Model Evaluation Results

To evaluate the performance of the trained model, several metrics were used, including Inception Score (IS), Fréchet Inception Distance (FID), and a user study for subjective quality assessment. The evaluation results are summarized in Table 2.

Table 2. Model Evaluation Metrics Comparison		
Metric	Proposed Workflow	Baseline Model
Inception Score (IS)	3.25	2.85
Fréchet Inception Distance (FID)	18.5	22.3

The Inception Score (IS) evaluates the quality and diversity of the generated images. A higher IS indicates more realistic and varied images. The proposed workflow achieved an IS of 3.25, outperforming the baseline model, which scored 2.85. This improvement is attributed to the effective conditioning mechanism in the Conditional Generative Adversarial Network (cGAN) architecture, which enhances the diversity of the generated images. The Fréchet Inception Distance (FID) measures the similarity between the generated images and the real images from the test set. A lower FID indicates that the generated images are more similar

to real images. The proposed workflow achieved an FID of 18.5, significantly better than the baseline model's FID of 22.3, demonstrating improved realism and consistency with the conditional input. A user study was conducted to evaluate the perceived quality and realism of the generated images. Participants rated the images on a scale of 1 to 5, and the results are shown in Figure 2. The proposed workflow consistently received higher ratings compared to the baseline model, validating the effectiveness of the automated model fine-tuning approach.

4.4 Model Deployment Results

After achieving satisfactory model performance, the best model was deployed to a production environment using AWS SageMaker's managed endpoints. The deployment process involved creating a model artifact and deploying it to an endpoint configured with a T2 instance for cost-effective scalability. The deployed endpoint provides a REST API for real-time image generation, enabling seamless integration with web and mobile applications. To evaluate the deployment's performance, a series of sample requests were sent to the endpoint. The generated images were visually inspected for quality and consistency with the input conditions. The endpoint demonstrated high availability and low latency, efficiently handling a high number of requests without performance degradation. The auto-scaling feature was enabled to dynamically adjust the instance count based on the incoming traffic, ensuring a responsive user experience under varying loads. Moreover, model versioning was implemented to track different versions of the model, enabling continuous model improvement and rollback capabilities if needed. This approach enhances the maintainability and scalability of the deployment pipeline. In conclusion, the experimental results validate the effectiveness of the proposed workflow for automated model fine-tuning and deployment using AWS SageMaker. The optimized hyperparameters and advanced cGAN architecture significantly improved the image quality and diversity, outperforming the baseline model. The deployment strategy ensured high availability and scalability, making the model suitable for real-world applications.

5. Discussion

This section provides an in-depth analysis of the experimental results, highlighting the effectiveness, scalability, and limitations of the proposed workflow for automated model fine-tuning and deployment using AWS SageMaker. Additionally, potential avenues for future work are discussed to enhance the workflow's performance and applicability.

5.1 Performance Analysis

The experimental results demonstrate that the proposed workflow significantly enhances the performance of the image generation model. By leveraging AWS SageMaker's automatic model tuning, the workflow efficiently explores the hyperparameter space, resulting in an optimized model configuration. The improved hyperparameters contribute to superior performance, as evidenced by the higher Inception Score (IS) and lower Fréchet Inception Distance (FID) compared to the baseline model. The Inception Score (IS) of 3.25 indicates that the generated images are not only realistic but also diverse, reflecting the model's ability to capture complex patterns and features from the training dataset. Similarly, the lower FID score of 18.5 demonstrates the model's capability to generate images that are perceptually closer to real images. These quantitative metrics are further validated by the user study results, where participants consistently rated the images generated by the proposed workflow higher in terms of quality and realism. The performance improvements can be attributed to several factors, including the use of Conditional Generative Adversarial Networks (cGANs), which enhance the model's controllability and diversity. Additionally, the Adam optimizer's adaptive learning rate and momentum properties contribute to faster convergence and improved generalization. The overall performance analysis confirms that the proposed workflow effectively balances quality, diversity, and computational efficiency, making it a robust solution for image generation tasks.

5.2 Scalability and Flexibility

One of the key strengths of the proposed workflow is its scalability and flexibility. AWS SageMaker's managed services enable the workflow to handle large datasets and complex models without the need for extensive infrastructure management. The use of P3 instances with Tesla V100 GPUs ensures efficient computation, reducing training time and enabling rapid experimentation.

The deployment strategy leverages SageMaker endpoints, which provide a seamless transition from model training to deployment. These endpoints offer auto-scaling capabilities, allowing the model to dynamically adjust its computational resources based on traffic demands. This ensures high availability and low latency, even under fluctuating workloads. Moreover, the workflow supports model versioning, enabling continuous model improvement and easy rollback capabilities in case of performance degradation. This versioning capability enhances maintainability and facilitates rapid deployment of updated models. The flexibility of the workflow is further demonstrated by its compatibility with various image generation tasks, including conditional and unconditional image synthesis. This adaptability makes it suitable for a wide range of applications, from creative content generation to computer vision tasks in industrial settings.

5.3 Limitations and Future Work

Despite the advantages offered by the proposed workflow, there are several limitations to consider:

1. **Computational Cost:** The hyperparameter optimization process, while effective in enhancing model performance, can be computationally expensive, especially when exploring large hyperparameter spaces. This could result in increased training costs, particularly when using high-performance GPUs.
2. **Model Complexity:** The complexity of the Conditional Generative Adversarial Network (cGAN) architecture can lead to longer training times and higher resource requirements. This may limit the applicability of the workflow for smaller organizations with budget constraints.
3. **Data Privacy:** The use of cloud-based services raises concerns about data privacy and security, especially when dealing with sensitive or proprietary datasets. Ensuring compliance with data protection regulations such as GDPR and CCPA remains a challenge.

To address these limitations, several directions for future work are proposed:

- **Enhanced Hyperparameter Optimization:** Future research could explore more advanced hyperparameter optimization techniques, such as evolutionary algorithms and reinforcement learning-based methods. These techniques could further improve the model's performance while reducing computational costs.
- **Multi-Task Learning:** The workflow could be extended to support multi-task learning, enabling the model to generate images for multiple tasks simultaneously. This would enhance the workflow's versatility and efficiency in scenarios requiring diverse outputs.
- **Federated Learning:** To address data privacy concerns, future work could investigate the integration of federated learning. This approach would enable decentralized model training on edge devices, ensuring data privacy while maintaining model performance.

6. Conclusion

In conclusion, this paper presents a scalable and efficient workflow for automated model fine-tuning and deployment using AWS SageMaker, specifically tailored for image generation tasks. The workflow leverages SageMaker's advanced hyperparameter optimization capabilities to automate the tuning process, ensuring optimal model configurations with minimal manual intervention. The experimental results clearly demonstrate the superior performance of the proposed workflow compared to traditional methods. The optimized model achieved a higher Inception Score and lower Fréchet Inception Distance, indicating improved image quality and diversity. The user study further validated the enhanced realism and perceptual quality of the generated images.

The workflow's scalability is evident from its ability to handle large datasets and complex models while maintaining high availability and low latency through managed endpoints and auto-scaling features. Additionally, the flexibility of the workflow allows for seamless adaptation to various image generation tasks, making it a valuable tool for researchers and practitioners alike. While the workflow shows great promise, challenges such as computational cost, model complexity, and data privacy need to be addressed. Future work exploring advanced hyperparameter optimization techniques, multi-task learning, and federated learning could further enhance the workflow's efficiency, versatility, and security. Overall, the proposed workflow not only advances the state of the art in automated model fine-tuning and deployment but also provides a practical and scalable solution for real-world image generation applications. It serves as a foundation for further research and development in the field of automated machine learning and cloud-based model deployment.

References

- [1] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672-2680).
- [2] Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- [3] Amazon Web Services. (2021). AWS SageMaker. Retrieved from <https://aws.amazon.com/sagemaker/>
- [4] Zhang, H., Goodfellow, I., Metz, L., & Odena, A. (2018). Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*.
- [5] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., & Hochreiter, S. (2017). GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems* (pp. 6626-6637).
- [6] <https://docs.aws.amazon.com/sagemaker-unified-studio/latest/userguide/bedrock-image-playground-generate-image.html>
- [7] <https://www.pluralsight.com/resources/blog/ai-and-data/deploying-genAI-amazon-sagemaker>
- [8] <https://docs.aws.amazon.com/sagemaker/latest/dg/studio-byoi-create.html>
- [9] <https://aws.amazon.com/blogs/machine-learning/automate-fine-tuning-of-llama-3-x-models-with-the-new-visual-designer-for-amazon-sagemaker-pipelines/>
- [10] <https://www.youtube.com/watch?v=2t4Qiq7hhJ8>
- [11] <https://docs.aws.amazon.com/sagemaker/latest/dg/canvas-fm-chat-fine-tune.html>
- [12] <https://aws.amazon.com/blogs/machine-learning/prepare-image-data-with-amazon-sagemaker-data-wrangler/>
- [13] <https://docs.aws.amazon.com/sagemaker/latest/dg/jumpstart-fine-tune.html>

- [14] <https://aws.amazon.com/blogs/machine-learning/generate-unique-images-by-fine-tuning-stable-diffusion-xl-with-amazon-sagemaker/>
- [15] <https://aws.amazon.com/blogs/machine-learning/architect-personalized-generative-ai-saas-applications-on-amazon-sagemaker/>

Appendices

Appendix A: Algorithm for Hyperparameter Optimization

```
def hyperparameter_optimization(model, dataset, hyperparameter_space, metric, objective):  
    # Define the hyperparameter tuning job  
    tuner = HyperparameterTuner(  
        model,  
        objective_metric_name=metric,  
        hyperparameter_ranges=hyperparameter_space,  
        objective_type=objective,  
        max_jobs=10,  
        max_parallel_jobs=2  
    )  
  
    # Start the hyperparameter tuning job  
    tuner.fit({'training': dataset})  
  
    # Get the best hyperparameters  
    best_hyperparameters = tuner.best_hyperparameters()  
  
    return best_hyperparameters
```

Appendix B: Model Training Code

```
import sagemaker  
from sagemaker.pytorch import PyTorch  
  
# Define the model  
model = PyTorch(  
    entry_point='train.py',  
    role=sagemaker.get_execution_role(),  
    framework_version='1.8.1',  
    py_version='py36',  
    instance_count=1,  
    instance_type='ml.p3.2xlarge',  
    hyperparameters={  
        'learning_rate': 0.0002,  
        'batch_size': 64,  
        'epochs': 100,  
        'noise_dim': 100,  
        'condition_dim': 10,  
        'optimizer': 'adam'  
    }  
)  
  
# Train the model  
model.fit({'training': 's3://path/to/dataset'})
```

Appendix C: Model Deployment Code

```
import sagemaker  
  
# Create a model  
model = sagemaker.Model(  
    model_data='s3://path/to/model.tar.gz',  
    role=sagemaker.get_execution_role(),
```



```
framework_version='1.8.1',  
py_version='py36',  
entry_point='inference.py'  
)  
  
# Deploy the model  
predictor = model.deploy(  
    initial_instance_count=1,  
    instance_type='ml.t2.medium'  
)  
  
# Test the endpoint  
response = predictor.predict({'noise': random_noise, 'condition': condition_vector})  
generated_image = response['image']
```