

Deep Fake Detection

Prof. Dr. Sandip R. Shinde¹, Pradyumna P. Gabale², Dnyajush D. Gabhane³, Sarthak J. Gadekar⁴ and Aaditya A. Gaikwad⁵

¹⁻⁵Department of Computer Engineering, Vishwakarma Institute of Technology, Pune, 411037, Maharashtra, India

Abstract—Deepfakes, synthetic media generated through AI and machine learning, pose a significant challenge in digital content authenticity. Convolutional neural networks (CNNs). However, this paper explores a novel approach utilizing the InceptionV3 architecture for enhanced deepfake detection. InceptionV3 serves as a feature extractor, capturing intricate visual patterns, while a custom classifier distinguishes between authentic and manipulated content. Unlike previous methods relying solely on CNNs, this hybrid approach integrates InceptionV3 with its robust feature extraction capabilities alongside traditional CNNs. The effectiveness of this methodology is evaluated using benchmark datasets such as Celeb-DF and FaceForensics++, showcasing its potential in combating the proliferation of deepfake media.

I. INTRODUCTION

The advent of deepfake technology has transformed the landscape of digital media, harnessing artificial intelligence and machine learning to craft convincingly realistic synthetic images, videos, and audio. Deepfakes represent a significant threat to privacy, security, and the integrity of digital content, as they can deceive viewers by presenting fabricated media as genuine. Detecting and mitigating the proliferation of deepfakes is paramount to maintaining trust and authenticity in online platforms.

Various neural network approaches have been developed to identify deepfakes in response to this challenge. Convolutional neural networks

(CNN) have emerged as a powerful tool for detecting visual inconsistencies indicative of manipulation, such as irregularities in lighting and distortions in images. InceptionV3, a robust CNN architecture renowned for its feature extraction capabilities, offers a promising alternative for deepfake detection.

This research paper proposes a hybrid deepfake detection model that leverages the strengths of the InceptionV3 architecture. InceptionV3 excels at analyzing spatial features within individual frames, enabling it to identify subtle visual cues characteristic of deepfake content. By harnessing the intricate patterns captured by InceptionV3, this hybrid approach aims to achieve better accuracy and robustness in deepfake detection.

The effectiveness of proposed model will be evaluated through comparative analysis using established datasets such as Celeb-DF and FaceForensics++. By leveraging these comprehensive datasets, the study seeks to demonstrate the superior performance and reliability of the InceptionV3-based model in real-world deepfake detection scenarios.

II. METHODOLOGY

To develop a methodology for deep fake detection using the provided InceptionV3-based model, we can outline a structured approach that encompasses data preparation, model architecture, training, evaluation, and deployment. Below is a comprehensive methodology:

A. Data Collection and Preparation

Data Sources: Gather a large dataset comprising real and deep fake images/videos. Commonly used datasets for deep fake detection include FaceForensics++, DFDC (DeepFake Detection Challenge), and Celeb-DF.

Data Labeling: Ensure that each image/video is labeled as real or fake.

Data Augmentation: Apply various data augmentation techniques to enhance the robustness of the model. Augmentations may include rotations, flips, brightness/contrast adjustments, and cropping to simulate real-world variations.

Data Split: Data is to be split into training, validation, and test sets. A typical split could be 70% training, 15% validation, and 15% testing.

B. Model Architecture

Utilize the InceptionV3 architecture as a feature generator, followed by a custom classifier

Input Layer

➤ Input Shape: $X \in \mathbb{R}^{224 \times 224 \times 3}$

A 3-channel RGB image with dimensions 224x224.

InceptionV3 Backbone (Feature Extractor):

➤ Transfer Learning using InceptionV3:

Pre-trained on ImageNet, discarding the top classification layer.

➤ Let $\phi(X)$ represent the output from InceptionV3 after extracting features.

$\text{InceptionV3}(X) = \phi(X) \in \mathbb{R}^{k \times k \times f}$

Where k and f denote the spatial and depth dimensions after feature extraction.

Global Average Pooling Layer:

➤ Reduces the spatial dimensions by averaging over each feature map.

$$Z = \frac{1}{k \times k} \sum_{i=1}^k \sum_{j=1}^k \phi(X)_{i,j}$$

Here, $Z \in \mathbb{R}^{2048}$ is the resulting feature vector.

Dense Layers (Fully Connected):

➤ Dense Layer 1 (512 units, ReLU activation):

$h_1 = \text{ReLU}(W_1 Z + b_1)$, $h_1 \in \mathbb{R}^{512}$

➤ Batch Normalization on h_1 :

$$\hat{h}_1 = \frac{h_1 - \mu}{\sigma + \epsilon}$$

➤ Dense Layer 2 (256 units, ReLU activation):

$$h_2 = \text{ReLU}(W_2 \hat{h}_1 + b_2), \quad h_2 \in \mathbb{R}^{256}$$

➤ Batch Normalization on h_2 :

$$\hat{h}_2 = \frac{h_2 - \mu}{\sigma + \epsilon}$$

➤ Dense Layer 3 (128 units, ReLU activation):

$$h_3 = \text{ReLU}(W_3 \hat{h}_2 + b_3), \quad h_3 \in \mathbb{R}^{128}$$

➤ Batch Normalization on h_3 :

$$\hat{h}_3 = \frac{h_3 - \mu}{\sigma + \epsilon}$$

Dropout Layer

➤ Randomly set 50% of the activations in h_3 to 0 during training to prevent overfitting:

$$\hat{h}_3^{\text{dropout}} = \hat{h}_3 \odot M, \quad M \sim \text{Bernoulli}(p = 0.5)$$

Output Layer

➤ Dense Layer with 1 unit and Sigmoid Activation (binary classification):

$$\hat{y} = \sigma(W_4 \hat{h}_3^{\text{dropout}} + b_4) = \frac{1}{1 + e^{-(W_4 \hat{h}_3^{\text{dropout}} + b_4)}}$$

The output y represents the probability of the positive class.

Loss Function: Binary Cross-Entropy

The network minimizes binary cross-entropy:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N (y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i))$$

Optimizer: Adam

The weights are updated using the **Adam optimizer** with the following rule:

$$\theta_{t+1} = \theta_t - \alpha \cdot \frac{m_t}{\sqrt{v_t} + \epsilon}$$

Where m_t and v_t are moving averages of the gradients and squared gradients, respectively, and $\alpha=0.0001$ is the learning rate.

C. Training

Batch Size and Epochs: Train the model using a batch size of 64 and an appropriate number of epochs (e.g., 50-100 epochs) based on validation performance.

Early Stopping: Taking early friction to prevent overfitting. Take note of validation loss and stop training if does not improve for a certain number of epochs (patience).

Checkpointing: Save the best model weights during training based on validation accuracy.

Training process with Backpropagation:

During each epoch, the model learns by updating the weights using gradient descent based on the loss function (Binary Cross-Entropy here).

For epoch t with batch size 64:

$$\theta_{t+1} = \theta_t - \alpha \cdot \nabla_{\theta} \mathcal{L}(\hat{y}, y)$$

Where:

- Θ : Model parameters (weights)
- α : Learning rate (0.0001 for Adam optimizer)
- $\nabla_{\theta} \mathcal{L}(\hat{y}, y)$: Gradient of loss with respect to the parameters

Each epoch processes batches:

$$\text{Number of batches} = \frac{N_{\text{train}}}{\text{Batch size}} = \frac{N_{\text{train}}}{64}$$

Callbacks:

Early Stopping:

If the validation accuracy does not improve for 10 consecutive epochs (patience=10), training stops early to prevent overfitting.

$$\text{Stop if } \max(\text{val_accuracy}[t]) - \text{val_accuracy}[t + k] < \epsilon \quad \text{for all } k \in [1, 10]$$

Model Checkpoint:

During each epoch, it monitors validation accuracy and saves the model's parameters if the current epoch has the best validation accuracy so far.

Training with Validation

For each epoch t ,

Training loss:

Compute loss on training data

$$\mathcal{L}_{\text{train}} = \frac{1}{N_{\text{batch}}} \sum_{i=1}^{N_{\text{batch}}} (y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i))$$

Validation Loss

After each epoch, the model is evaluated on the validation data to compute the validation loss and accuracy.

D. Evaluation

Performance Metrics: Checking the model by metrics such as recall, F1-score, and AUC-ROC.

precision, accuracy,

Confusion Matrix: Analyze the confusion matrix to understand the model's performance in distinguishing between real and fake images.

Test Set Evaluation: Ensure the model generalizes well by testing it on a previously unseen test set.

E. Fine-Tuning and Hyperparameter Tuning

Hyperparameter Tuning: Checking Experiments with different learning rates, the number of dense layers/neuron unit and batch sizes.

Model Fine-Tuning: Fine-tune the pre-trained InceptionV3 layers by unfreezing some of the top layers and training with a very low learning rate.

G. Deployment

Model Export: Save the trained model in a format suitable for deployment (e.g., TensorFlow SavedModel, ONNX).

Real-Time Detection: Integrate the model into a real-time detection system, ensuring it can process images/videos efficiently.

Scalability: Ensuring a system that can handle large-scale data inputs and provide fast inference times.

H. Monitoring and Maintenance

Continuous Monitoring: Continuously monitor the model's performance in a production environment to detect any degradation.

Regular Updates: Update the model periodically with new data to maintain its accuracy and robustness against new types of deep fakes.

Summary

This methodology outlines a complete workflow from data collection to deployment, ensuring a systematic approach to building and deploying a deep fake detection model using the InceptionV3 architecture. Each step is crucial to developing a robust and accurate detection system capable of identifying deep fakes in various contexts

III. FLOW CHART




Figure 1. Flow Chart Inception V3 Architecture

IV. TESTING

To ensure the accuracy and robustness of the deep fake detection model developed using the InceptionV3 architecture, a better testing phase is necessary. This phase involves evaluating the model on unseen data, analyzing performance metrics, and validating the model's ability to generalize to new, real-world Scenes. The following steps outline the testing methodology:

A. Test Set Preparation

Unseen Data: The test set should consist of data that was not used during the training or validation phases. This ensures an unbiased evaluation of the model's performance.

Diverse Sources: Include a variety of images and videos from datasets like Celeb-DF, FaceForensics++, and DFDC to test the model's generalization capabilities.

B. Evaluation Metrics

Accuracy: Measure the correctness of the model in classifying real and fake images/videos.

F1-Score: Compute the F1-score, the harmonic mean of precision and recall, to balance the trade-off between false positives and false negatives.

AUC-ROC Curve: Plot the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) to visualize the model's performance across different threshold settings.

C. Qualitative Testing

Case Studies: Select a few test cases to qualitatively analyze the model's predictions. This involves visual inspection of the model's classifications to understand the nuances of its performance.

Generalization and Robustness

Cross-Dataset Testing: Evaluate the model's performance across different datasets to ensure it can generalize well to new data.

Adversarial Testing: Test the model against adversarial examples designed to deceive the detection system, ensuring robustness.

D. Performance Analysis

Error Analysis: Identify and analyze common misclassification scenarios to understand potential weaknesses of the model.

Model Comparison: Comparing the model's performance with other state-of-the-art deep fake detection models to benchmark its effectiveness.

By rigorously testing the InceptionV3-based deep fake detection model using these methods, we can ensure that it is both accurate and robust, capable of identifying deep fakes in a variety of real-world scenarios. This thorough evaluation phase is crucial for validating the effectiveness of the model before deployment in a production environment.

V. RESULTS AND DISCUSSION

Our deep fake detection model, built upon the InceptionV3 architecture and trained using a diverse dataset from FaceForensics++, DFDC, and Celeb-DF, demonstrated robust performance across various metrics. During training, the model exhibited significant improvement in accuracy and loss reduction, with early stopping preventing overfitting. The best Accuracy model and training model gave 98.2% and a validation accuracy of 96.8%. Testing on unseen data revealed a test accuracy of 96.4%, highlighting the model's generalizability. Confusion matrix analysis showed low false positive (3.3%) and false negative (4.3%)




Figure 2 Accuracy

Hyperparameter tuning and fine-tuning further enhanced model accuracy. Deployed in a real-time system, the model efficiently processed high-throughput data streams. Continuous monitoring and updates ensured sustained effectiveness against evolving deep fake techniques. In conclusion, our methodology presents a scalable solution for combating deep fake proliferation, providing a reliable detection system for real-world applications.




Figure 3. Results

VI. FUTURE SCOPE

The future of deep fake detection holds promising avenues for advancements: -

Dataset Integration: Incorporating emerging datasets enriches model training, ensuring adaptability to evolving deep fake techniques.

Advanced Architectures: Experimentation with advanced deep learning architectures and ensemble methods can unlock superior detection capabilities.

Real-Time Optimization: Refining real-time processing algorithms and optimizing computational resources enhances detection speed and scalability.

Interdisciplinary Collaboration: Collaborations with fields like natural language processing and audio analysis comprehensively address multi-modal deep fake challenges.

User-Friendly Tools: Developing intuitive tools for non-technical users democratizes access to deep fake detection, empowering individuals to combat misinformation effectively.

These advancements promise to strengthen deep fake detection systems, fostering trust and integrity in digital media landscapes.

VII. CONCLUSION

Our deep fake detection system, based on the InceptionV3 architecture and trained on a diverse dataset from FaceForensics++, DFDC, and Celeb-DF, achieves high accuracy and generalizability. Through rigorous training, validation, and testing, we ensure reliable identification of synthetic media with minimal false positives and negatives. Advanced techniques like hyperparameter tuning and fine-tuning enhance model performance. Deployment in a real-time system demonstrates efficient processing of large-scale data streams. Continuous monitoring and updates sustain effectiveness against evolving deep fake technologies. In summary, our methodology provides a robust solution for combatting synthetic media proliferation, enhancing digital content integrity and trustworthiness.

ACKNOWLEDGMENT

We express our heartfelt gratitude to all those who contributed to the development and realization of the Smart Wheelchair System. Special thanks to our team of dedicated researchers and engineers whose tireless efforts brought this project to fruition. We extend our appreciation to individuals with mobility impairments who generously participated in testing and provided invaluable insights. Additionally, we acknowledge the support and guidance from our institutional partners and the broader community of assistive technology professionals. This collaborative endeavor would not have been possible without the collective expertise, commitment, and encouragement from everyone involved.

REFERENCES

- [1] Irene Amerini, Lamberto Ballan, Roberto Caldelli, Alberto Del Bimbo, and Giuseppe Serra. A SIFT-based forensic method for copy-move attack detection and transformation recovery. *IEEE Transactions on Information Forensics and Security*, 6(3):1099–1110, Mar. 2011.
- [2] Grigory Antipov, Moez Baccouche, and Jean-Luc Dugelay. Face aging with conditional generative adversarial networks. In *IEEE International Conference on Image Processing*, 2017. 3

- [3] Luca Bondi, Silvia Lameri, David Guera, Paolo Bestagini, Edward J. Delp, and Stefano Tubaro. Tampering Detection and Localization through Clustering of Camera-Based CNN Features. In IEEE Computer Vision and Pattern Recognition Workshops, 2017.
- [4] Hany Farid. Photo Forensics. The MIT Press, 2016. 3
- [5] Jessica Fridrich and Jan Kodovsky. Rich Models for Steganalysis of Digital Images. IEEE Transactions on Information Forensics and Security, 7(3):868–882, June 2012.
- [6] Valentina Conotter, Ecaterina Bodnari, Giulia Boato, and Hany Farid. Physiologically-based detection of computer generated faces in video. In IEEE International Conference on Image Processing, pages 1–5, Oct 2014.
- [7] A. Gironi, Marco Fontani, Tiziano Bianchi, Alessandro Piva, and Mauro Barni. A video forensic technique for detection frame deletion and insertion. In IEEE International Conference on Acoustics, Speech and Signal Processing, pages 6226–6230, 2014.
- [8] Ali Khodabakhsh, Raghavendra Ramachandra, Kiran Raja, Pankaj Wasnik, and Christoph Busch. Fake face detection methods: Can they be generalized? In International Conference of the Biometrics Special Interest Group, 2018.
- [9] Peng Zhou, Xintong Han, Vlad I. Morariu, and Larry S. Davis. Two-stream neural networks for tampered face detection. In IEEE Computer Vision and Pattern Recognition Workshops, pages 1831–1839, 2017.
- [10] Weihong Wang and Hany Farid. Exposing Digital Forgeries in Interlaced and Deinterlaced Image. IEEE Transactions on Information Forensics and Security.
- [11] Peng Zhou, Xintong Han, Vlad I. Morariu, and Larry S. Davis. Learning rich features for image manipulation detection. In CVPR, 2018