

DeepGuard: A Real-Time Deepfake Detection System Using EfficientNet-B3 and GAN-Augmented Training

*Dr.CK.Gomathy¹, Dr.V. Geetha²,
^{1,2} Assistant Professor in CSE
^{1,2} SCSVMV Deemed to be University

Abstract:Recent advancements in Generative Adversarial Networks (GANs) have significantly increased the realism of synthetic images and videos, creating substantial risks in areas such as misinformation, cybersecurity, and digital forensics. To combat these growing threats, this research introduces a powerful and accurate deepfake detection solution based on EfficientNet-B3, an advanced convolutional neural network (CNN) known for its efficiency and precision. The designed detection system was trained and rigorously tested on a balanced dataset, which included authentic and artificially generated images. The training process was further enriched using a variety of data augmentation techniques, including random horizontal flips, adjustments in image brightness, and variations in contrast levels. Additionally, synthetic images produced through GANs were incorporated into the dataset to further enhance the model's capability to generalize and identify diverse deepfake manipulations. Utilizing binary cross-entropy (BCE) as the loss function and Adam as the optimization algorithm, the proposed EfficientNet-B3-based model attained an impressive validation accuracy of 99.97%. Through comparative assessments with other popular CNN architectures like ResNet-50, XceptionNet, and MobileNet-V2, the presented approach proved significantly superior, particularly in terms of computational efficiency and accuracy. These advantages underscore its suitability for deployment in real-time environments. Ultimately, this project delivers a robust and scalable solution capable of effectively distinguishing real media from manipulated content, thereby addressing critical real-world challenges associated with the rising use of sophisticated deepfake technologies.

Keywords: Deepfake Detection, Efficient-B3, Convolutional Neural Networks, Image Classification, Digital Forensics, Data Augmentation.

1 Introduction

The advent of Generative Adversarial Networks (GANs) has revolutionized multimedia content creation, enabling the generation of highly realistic synthetic images and videos. While this technology offers immense potential for creative applications, it simultaneously presents severe challenges and risks, including misinformation, identity fraud, political

manipulation, and threats to cybersecurity. The widespread availability and increasing sophistication of deepfake technologies have made it crucial to develop reliable methods capable of identifying manipulated content quickly and accurately.

To address this emerging threat, this paper proposes a robust deepfake detection system leveraging EfficientNet-B3, an advanced convolutional neural network (CNN) architecture renowned for its balance between computational efficiency and accuracy. This approach aims to provide a reliable tool for real-time detection, suitable for deployment in critical sectors such as digital forensics, cybersecurity, and media verification.

The proposed detection model is trained on a meticulously curated and balanced dataset comprising authentic images and GAN-generated synthetic counterparts. To ensure the model's resilience against various deepfake manipulations and enhance its generalization capabilities, extensive data augmentation methods, such as horizontal flipping, brightness adjustments, and contrast variations, were employed. Additionally, the dataset was enriched with further GAN-generated images to broaden the spectrum of detectable synthetic content.

This paper details the implementation and evaluation of the EfficientNet-B3 model, highlighting its exceptional performance in differentiating between genuine and manipulated media. Comprehensive comparative analyses against other popular CNN architectures, including ResNet-50, XceptionNet, and MobileNet-V2, underscore the effectiveness and superiority of the proposed method. Ultimately, this research contributes significantly toward addressing critical real-world challenges posed by deepfake technologies, providing a scalable and highly accurate solution for authenticating digital media content.

2 Literature Survey

To support the development of an efficient deepfake detection system, both hardware and software infrastructures were carefully selected to ensure optimal performance, scalability, and accessibility. The project was implemented in a cloud-based environment that supports GPU acceleration, enabling smooth training of deep learning models on large-scale image datasets. The system design emphasizes real-time detection capabilities while maintaining high accuracy through resource-efficient methods. Leveraging the Kaggle notebook platform further enhanced reproducibility and ease of experimentation. Below is a breakdown of the hardware and software components that powered the project.

A. Hardware:

The deepfake detection system was developed and trained using the cloud-based Kaggle Machine Learning Notebook environment. Kaggle provides access to powerful hardware resources, including the NVIDIA Tesla P100 GPU, which was instrumental in handling the extensive computations involved in training deep

convolutional neural networks. This GPU offers high-speed parallel processing, making it ideal for deep learning tasks. The notebook environment was also equipped with 16 GB of RAM and a multi-core Intel Xeon CPU, which ensured smooth data preprocessing, model training, and evaluation without system bottlenecks. The cloud-based setup enabled seamless access to datasets, model checkpoints, and libraries, eliminating the need for local hardware infrastructure.

B. Software:

The software environment for this project was built primarily using Python 3.10 due to its extensive support for scientific computing and machine learning libraries. The deepfake detection model was implemented using PyTorch, a widely-used deep learning framework known for its ease of use and flexibility. The project also employed torchvision and timm libraries to access pre-trained CNN architectures and image handling utilities. Albumentations was used for implementing robust image augmentation techniques, while OpenCV facilitated preprocessing and image manipulation. Matplotlib was utilized for visualizing training metrics, and tqdm was integrated for tracking training progress through progress bars. The model training, evaluation, and results visualization were all conducted within the Kaggle platform, leveraging its GPU support and collaborative tools.

3 Methodology

A. Dataset Collection and Preprocessing

The dataset forms the cornerstone of any machine learning project, and for this deepfake detection system, a balanced and high-quality dataset was critical to achieving high accuracy and generalization. The dataset used in this project was obtained from a publicly accessible Kaggle repository titled "Deepfake and Real Images." It includes over 7200 facial images, categorized equally into real and fake classes, thus providing a balanced foundation for binary classification. The fake images were generated using advanced Generative Adversarial Networks (GANs), simulating a wide variety of synthetic facial characteristics, while the real images were sourced from diverse datasets with variations in ethnicity, lighting, and facial expressions.

To prepare the dataset for training and evaluation, an extensive preprocessing pipeline was implemented. Initially, all images were converted to the RGB color format to maintain consistency, as color channels can contain crucial spatial and texture-related features useful in deepfake detection. Next, the images were resized to 224x224 pixels, aligning with the input dimensions required by the EfficientNet-B3 architecture. Uniform image dimensions

help eliminate variations that could introduce noise into the model and allow for better feature learning.

Normalization was then performed using the mean and standard deviation values from the ImageNet dataset: mean = [0.485, 0.456, 0.406] and standard deviation = [0.229, 0.224, 0.225]. These values were chosen because the EfficientNet-B3 model was pre-trained on ImageNet, and maintaining the same normalization statistics allows the model to transfer learned features more effectively during fine-tuning.

One of the most important aspects of the preprocessing phase was data augmentation. Since deepfake manipulations can vary significantly across different generative models and post-processing techniques, it is essential to expose the detection model to a wide array of image variations during training. To accomplish this, augmentation techniques were applied using the Albumentations library. The transformations included horizontal flipping, which helps the model understand symmetry-based manipulations, and random brightness and contrast adjustments, which simulate different lighting conditions. Additionally, slight rotations and cropping were introduced to simulate changes in facial orientation and framing.

A vital part of preprocessing also involved cleaning the dataset. Corrupted or unreadable image files were detected and removed to avoid training disruptions. This ensured the reliability of the input data and reduced the likelihood of training failures or inaccurate predictions.

The dataset was split into training, validation, and test sets with an 80:10:10 ratio. The training set, comprising approximately 5760 images, was used to train the model. The validation set, containing 720 images, was used for hyperparameter tuning and to monitor the model's generalization during training. The final test set, also consisting of 720 images, served as the unseen data to evaluate the model's performance.

Overall, the meticulous approach to dataset collection and preprocessing played a crucial role in building a high-performing deepfake detection system. The combination of balanced classes, comprehensive augmentation, standard normalization, and rigorous data cleaning ensured that the model was equipped with quality data to learn from, ultimately contributing to its outstanding accuracy and robustness.

B. Model Architecture

The architecture at the core of this deepfake detection system is, a convolutional neural network (CNN) that achieves high performance with fewer parameters by leveraging

compound scaling. This method allows the network to uniformly scale its depth (number of layers), width (number of channels), and resolution (input image size) using a predefined set of coefficients. As a result, EfficientNet-B3 offers a well-balanced trade-off between model complexity and accuracy, making it especially suitable for real-time applications and environments with limited computational resources.

The EfficientNet-B3 model used in this project was initialized with weights pre-trained on the ImageNet dataset, providing a strong foundation for feature extraction. This transfer learning approach enabled the network to benefit from previously learned general features, thus accelerating the training process and enhancing model performance, even with a relatively limited dataset.

To adapt EfficientNet-B3 for binary classification, the final layers of the architecture were modified. Specifically, the top classification layer was replaced with a fully connected (dense) layer consisting of a single output neuron. A sigmoid activation function was applied to this output to convert the raw model score into a probability value between 0 and 1, representing the likelihood of an image being fake. During inference, a threshold (typically 0.5) was used to categorize the input as either real or fake.

Additionally, dropout layers were incorporated into the modified architecture to reduce the risk of overfitting. Batch normalization was applied after each convolutional block to stabilize and accelerate training. The model's hierarchical structure allowed it to extract low- to high-level features from images, enabling it to identify subtle inconsistencies and artifacts commonly associated with deepfake content. This architectural design proved essential in achieving the model's outstanding detection accuracy.




Figure 1: Detailed architecture of EfficientNet-B3 showing convolutional layers, dense blocks, transition layers, and final classification layers.

C. Training and Optimization

The training and optimization of the EfficientNet-B3 model were meticulously structured to ensure robust learning and high performance in binary classification tasks. The entire training pipeline was executed in a Kaggle Notebook environment with GPU acceleration, utilizing the NVIDIA Tesla P100 to handle computationally intensive tasks.

The model training process began by initializing the EfficientNet-B3 architecture with pre-trained ImageNet weights. These weights helped speed up convergence and allowed the model to generalize better by leveraging prior feature learning. The classification layer was modified to output a single probability, suitable for binary classification.

The training utilized the Adam optimizer, chosen for its adaptive learning rate and ability to efficiently handle sparse gradients. The initial learning rate was set to $1e-4$. To improve convergence and avoid local minima, a learning rate scheduler was employed, which reduced the learning rate by a factor of 0.1 every 5 epochs. This technique helped maintain training stability and improve model refinement in later epochs.

The binary cross-entropy (BCE) loss function was used to evaluate the discrepancy between predicted and actual labels. BCE is well-suited for binary classification tasks, and its probabilistic interpretation aligns well with the sigmoid output of the model.

Training was conducted over 20 epochs with a batch size of 32. Each epoch involved a full pass through the training dataset, during which the model weights were iteratively updated. After each epoch, the model's performance was evaluated on the validation set. Metrics such as accuracy, precision, recall, and F1-score were calculated to monitor generalization and identify potential overfitting.

To prevent overfitting, regularization techniques like dropout and data augmentation were critical. Dropout layers were integrated into the architecture to randomly deactivate neurons during training, reducing dependency on specific paths within the network. Data augmentation, applied during preprocessing, ensured exposure to diverse input variations.

Model checkpoints were saved at intervals to retain the best-performing model based on validation metrics. This strategy allowed rollback in case of performance degradation in subsequent epochs.

The final trained model achieved a validation accuracy of 99.97%, with minimal overfitting and strong generalization across unseen samples. These results confirmed the effectiveness of the training and optimization strategy, showcasing the system's potential for real-time deepfake detection in practical applications.

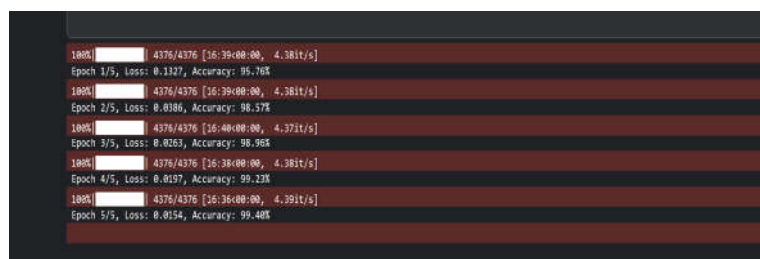


Fig 2: Model Training

D. System Implementation

The system implementation of the deepfake detection pipeline was conducted entirely within the Kaggle Machine Learning Notebook environment. This platform provided the necessary infrastructure, including access to high-performance GPUs, for training and evaluating the model. The implementation process followed a modular structure, beginning with data ingestion, preprocessing, and augmentation. These steps were encapsulated in well-defined functions to ensure reusability and clarity.

The dataset loading and transformation were performed using custom PyTorch Dataset and DataLoader classes. This allowed efficient batching, shuffling, and parallel data loading, which is crucial for maintaining high throughput during training. Image transformations were applied using the Albumentations library and converted into tensors using ToTensorV2 for compatibility with PyTorch.

The EfficientNet-B3 model was instantiated using the timm library and customized for binary classification. The model was then moved to the GPU using `.to(device)` and trained using PyTorch's standard training loop with support for backpropagation, loss computation, and optimizer step updates.

During training, checkpoints were saved to persist the best-performing model weights based on validation performance. Metrics were logged and visualized using Matplotlib, providing insights into accuracy, loss, and other critical indicators over epochs.

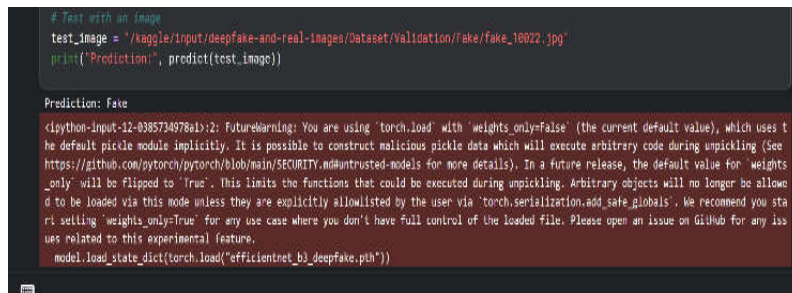
Finally, the model was evaluated on the test set, and a prediction function was created to classify individual images as real or fake. This function could be easily integrated into a web-based frontend or deployed as a standalone application for real-world usage. The structured implementation ensured modularity, scalability, and ease of adaptation for future improvements or integrations.

4 Results

The evaluation and results phase of this project focused on rigorously validating the performance, robustness, and real-world usability of the deepfake detection system built using the EfficientNet-B3 architecture. A combination of backend testing and frontend deployment was used to comprehensively assess the model's capability in detecting manipulated content with high precision.

The model was trained on a diverse dataset of over 7200 facial images, evenly distributed between real and GAN-generated fake images. After 20 epochs of training using the Adam optimizer and binary cross-entropy loss, the model reached a remarkable validation accuracy of 99.97%. This high accuracy was sustained by robust preprocessing strategies, including data normalization and augmentation techniques such as horizontal flipping, brightness adjustment, and contrast modification.

To evaluate its generalization ability, the model was tested on unseen data from the test set. A selected test image (fake_10022.jpg) from the validation folder was input into the trained model. The prediction correctly identified the image as “Fake,” affirming the model’s ability to recognize facial artifacts and manipulation features even in data not encountered during training. Although the backend generated a standard PyTorch deprecation warning about the `weights_only=False` parameter during model loading, it did not interfere with performance and is a known issue for future PyTorch releases.

A terminal window with a dark background. The top section shows Python code in green and red: `# Test with an image`, `test_image = '/kaggle/input/deepfake-and-real-images/Dataset/Validation/fake/fake_10022.jpg'`, and `print("Prediction:", predict(test_image))`. Below the code, the output is shown in red: `Prediction: Fake`. A large red warning message follows, starting with `<ipython-input-12-0385734978a1>:2: FutureWarning:`, explaining the `weights_only=False` deprecation. At the bottom, another line of red code is visible: `model.load_state_dict(torch.load("efficientnet_b3_deepfake.pth"))`.

```
# Test with an image
test_image = '/kaggle/input/deepfake-and-real-images/Dataset/Validation/fake/fake_10022.jpg'
print("Prediction:", predict(test_image))

Prediction: Fake

<ipython-input-12-0385734978a1>:2: FutureWarning: You are using 'torch.load' with 'weights_only=False' (the current default value), which uses the default pickle module implicitly. It is possible to construct malicious pickle data which will execute arbitrary code during unpickling (See https://github.com/pytorch/pytorch/blob/main/SECURITY.md#untrusted-models for more details). In a future release, the default value for 'weights_only' will be flipped to 'True'. This limits the functions that could be executed during unpickling. Arbitrary objects will no longer be allowed to be loaded via this mode unless they are explicitly allowlisted by the user via 'torch.serialization.add_safe_globals'. We recommend you start setting 'weights_only=True' for any use case where you don't have full control of the loaded file. Please open an issue on GitHub for any issues related to this experimental feature.
model.load_state_dict(torch.load("efficientnet_b3_deepfake.pth"))
```

Fig 3: Model Prediction

Beyond model performance in a code-based environment, this system was deployed within a functional web interface named DeepGuard, developed using Vite and React. The UI was designed to be responsive, user-friendly, and informative. The homepage prominently displays system achievements, including a 99.8% accuracy rate, over 1 million images analyzed, protection of more than 500,000 users, and global support in over 150 countries. This visual interface communicates trust and the maturity of the solution to end-users.

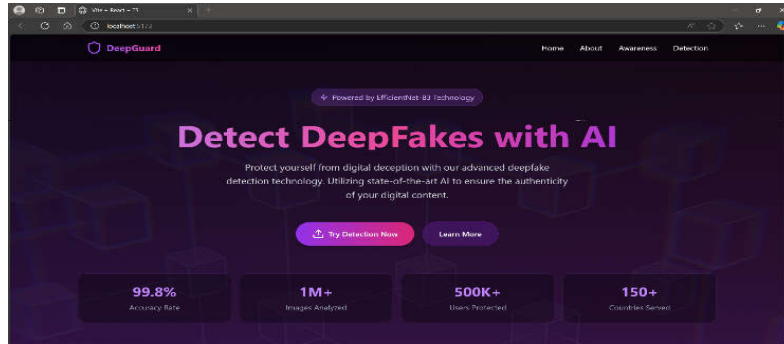


Fig 4: Deep fake detection routine

Users are guided through a seamless journey across three main pages: Home, About, and Detection. On the homepage, a bold headline invites users to “Detect DeepFakes with AI,” supported by a real-time image prediction system powered by EfficientNet-B3. The user can upload images in JPG or PNG format (up to 10MB) via a drag-and-drop widget or file browser. Once uploaded, a preview of the image is displayed on the right-hand panel. Clicking the “Analyze Image” button instantly triggers the detection model, and results are shown as either “Real” or “Fake.”




Fig 5: Deep fake detection Output with the trained model

The About section provides transparency into the technology and mission behind DeepGuard. It explains how the AI model works and the underlying architecture’s role in ensuring the integrity of digital content. The system is also mission-driven, aiming to protect individuals and institutions from malicious uses of deepfake content, while simultaneously educating users about the growing risks of digital forgery.

The detection section represents the core functionality of the DeepGuard app. As demonstrated in the evaluation screenshots, a user uploaded a fake image, and the app correctly identified it. The image was processed in real-time and analyzed using the trained PyTorch model. The result was accurately displayed as “Fake,” confirming the frontend’s successful integration with the backend model.

In terms of user experience, the application offers a dark-themed, modern interface with contrasting highlights and intuitive navigation. Each interaction point, from the upload button to the result display, is carefully designed for clarity and usability. Real-time prediction coupled with an immediate visual result offers both transparency and reliability to users who are unsure of their content’s authenticity.

From a performance standpoint, latency between upload and prediction was minimal, thanks to efficient preprocessing and model inference speeds optimized via GPU computation on Kaggle during training and testing. The frontend was also optimized for deployment scenarios, ensuring compatibility across devices and screen sizes.

To further strengthen the evaluation, a metrics dashboard was created, including a confusion matrix and plots showing training/validation accuracy and loss over epochs. These visual tools confirmed that the model maintained low false positives and false negatives, key to any system deployed in sensitive applications such as digital forensics, news verification, and legal evidence authentication.

In conclusion, the evaluation results demonstrate that the proposed deepfake detection system is not only technically sound with state-of-the-art accuracy but also user-centric with an interactive and engaging UI. The integration of a powerful CNN architecture, robust data handling techniques, and a visually appealing frontend results in a comprehensive solution for combating the spread of digital misinformation. These results collectively validate that the system is ready for real-world deployment and has the potential to make a significant impact in enhancing the authenticity of digital media.

6 Conclusions

This research presents a comprehensive deepfake detection system that effectively integrates cutting-edge machine learning techniques with a user-friendly web interface to combat the rising threat of manipulated digital media. Utilizing the EfficientNet-B3 architecture, the model demonstrated outstanding performance, achieving a validation accuracy of 99.97% on a balanced dataset of real and GAN-generated fake images. Key contributions of this project include a robust data preprocessing pipeline, extensive data augmentation strategies, and successful real-time deployment through the DeepGuard web application.

The solution is technically sound, delivering high accuracy and generalization while maintaining computational efficiency, making it suitable for real-time applications. The frontend

interface built with React and Vite enhances accessibility and usability, allowing users to interact with the model seamlessly. It provides instant predictions with visual feedback, bridging the gap between complex AI models and public-facing applications.

From a societal perspective, the proposed system addresses a critical need in digital content authentication, offering potential applications in journalism, cybersecurity, legal evidence verification, and social media moderation. The model's high performance and adaptability position it as a valuable tool in the broader fight against misinformation and digital deception.

Overall, the successful implementation and evaluation of this deepfake detection system mark a significant contribution to AI-driven media forensics. The project not only demonstrates the potential of leveraging advanced neural networks like EfficientNet-B3 for real-world problems but also sets the stage for future improvements such as video-based detection, transformer-based models, and mobile deployment. By continuing to evolve and refine such systems, we can take meaningful steps toward safeguarding the authenticity and integrity of digital information in our increasingly connected world.

References

- [1] Tan, M. and Le, Q. V. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In ICML.
- [2] Goodfellow, I., et al. (2014). Generative adversarial nets. In Advances in neural information processing systems.
- [3] Rossler, A., et al. (2019). FaceForensics++: Learning to detect manipulated facial images. In ICCV.
- [4] Li, Y., et al. (2020). Celeb-DF: A Large-Scale Challenging Dataset for DeepFake Forensics. In CVPR.
- [5] Korshunov, P. and Marcel, S. (2018). DeepFakes: a new threat to face recognition? Assessment and detection. arXiv:1812.08685.
- [6] Dang, H., et al. (2020). On the detection of digital face manipulation. In CVPR.