IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

DEEPFAKE DETECTION AND IMAGE RESTORATION SYSTEM

AN INTEGRATED APPROACH USING METADATA ANALYSIS, ELA, DEEP LEARNING, AND GAN-BASED DEBLURRING

[1] Bharti Sahu, [2]Rajashree Salunke , [3] Lalit Ushir, [4]Rohan Sawant , [5]Anisha Shinde

[1]Assistant Professor, [2]Student, [3] Student, [4] Student, [5] Student

Dr. D.Y. Patil Institute of Technology
Pimpri, Pune,India

Abstract: The proliferation of deepfake technology and the degraded image quality in digital media have necessitated comprehensive solutions for both detection and restoration. This paper presents an integrated AI-driven system that combines deepfake detection with GAN-based image restoration capabilities. Our approach uses EfficientNetB0 for deepfake classification and LPDGAN (Learned Primal-Dual Generative Adversarial Network) for image deblurring. The detection module uses transfer learning from ImageNet, achieving robust feature extraction on a dataset of 10,000 images (5,000 real, 5,000 fake). The restoration module employs 2,500 unsharp-sharp image pairs to train the LPDGAN architecture. The integrated system is deployed as a Flask-based web application, providing end-to-end processing from image upload to final classification through restoration. Experimental results demonstrate the effectiveness of pre-processing degraded images prior to detection, the system provides restored visual output for forensic analysis while successfully classifying images with high confidence. The modular architecture enables flexible deployment for real-world applications in digital forensics, journalism and cybersecurity.

Keywords-Deepfake detection, Image restoration, EfficientNetB0, LPDGAN, GAN-based deblurring, Deep learning.

I. INTRODUCTION

The rapid growth of digital content sharing on online platforms has created unprecedented challenges in verifying the authenticity and quality of visual media. Advanced editing tools and generative AI models have made it easier to create realistic image manipulation that can deceive both human observers and automated systems. Additionally, many forensic scenarios involve analyzing degraded, blurry, or low-quality images from surveillance footage, social media compression, or mobile captures that require restoration before reliable analysis.

Traditional approaches address deepfake detection and image restoration as separate problems, leading to fragmented workflows and suboptimal results. Deepfake detection methods have evolved from hand-crafted feature-based techniques to sophisticated deep learning architectures. However, these detectors often assume high-quality input images and perform poorly on coarse content. Similarly, image restoration techniques have advanced from simple filtering to GAN-based approaches, but are rarely integrated into forensic pipelines.

This research addresses both challenges by proposing an integrated system combining state-of-the-art deepfake detection with GAN-based image restoration. We recognize that real-world forensic applications require handling degraded evidence, which requires preprocessing before detection. Our system employs EfficientNetB0 for classification due to its optimal balance between accuracy and computational efficiency and LPDGAN for restoration due to its superior performance on motion blur and edge preservation.

The main contributions of this work include:

This paper introduces an integrated forensic system designed to detect deepfakes by combining advanced visual analysis with structural and metadata-based verification methods. The approach incorporates an EfficientNetB0-based image classifier alongside tools such as Error Level Analysis (ELA), EXIF data examination, and analysis of JPEG Huffman table structures. By fusing these visual and non-visual techniques, the system can uncover evidence of manipulation that might otherwise persist through common image compression and editing processes. Additionally, the framework features an object-focused restoration pipeline, utilizing YOLO for object detection and LPDGAN for image deblurring, to recover and enhance compromised forensic evidence. The entire solution is accessible through a web-based platform, allowing users to perform restoration, detection, and forensic analysis seamlessly within a single, scalable environment.

II.RELATED WORK RELATED WORK

A. Deepfake Detection Methods

Early deepfake detection methods relied on statistical analysis of hand-crafted features and image artifacts. The techniques examined anomalies in EXIF metadata, JPEG compression patterns, and color channel correlations. While computationally efficient, these approaches lacked robustness compared to sophisticated generation methods. The advent of deep learning revolutionized recognition capabilities. Convolutional neural networks enable automated feature learning while capturing subtle manipulation artifacts invisible to traditional methods. Transfer learning has proven particularly effective from large-scale datasets such as ImageNet, allowing models to take advantage of pre-learned visual representations.

Recent architectures focus on attention mechanisms and multi-level processing. Vision Transformers and hybrid CNN-Transformer models have shown promising results by capturing both local texture patterns and global semantic context. However, most research assumes high-quality input, with limited exploration of detection performance on degraded images.

B. EfficientNet Architecture

EfficientNet represents a breakthrough in neural architecture design through mixed scaling of network depth, width, and resolution. Unlike traditional approaches that arbitrarily scale one dimension, EfficientNet uses a theoretical compound coefficient to scale all three dimensions equally, achieving a better accuracy-efficiency trade-off. The base EfficientNetB0 model achieves state-of-the-art performance on ImageNet despite being much smaller than the ResNet and DenseNet variants.

To detect deepfakes, EfficientNet's mobile inverted bottleneck (MBConv) blocks enable efficient feature extraction from high-resolution facial images with squeeze-and-excitement optimization. The compound scaling of the architecture allows adaptation to different computational budgets while maintaining detection accuracy. Transfer learning from ImageNet provides robust initialization, which is important for scenarios with limited domain-specific training data.

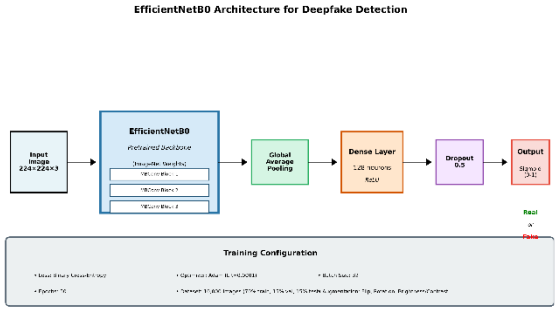

Fig.1. EfficientNetB0 Architecture

C. Image Deblurring Techniques

Image blurring has evolved from classical optimization-based methods to data-driven deep learning approaches. Traditional techniques model the blur as convolutions with point spread functions, using regularization and iterative optimization for restoration. These methods struggled with non-uniform blurring, noise, and computational efficiency. Deep learning revolutionized deblurring through end-to-end learning of complex image priors.

Generative Adversarial Networks introduced perceptual realism into restoration tasks. DeblurGAN demonstrated that adversarial training produces faster, more realistic outputs than pure L1/L2 reconstruction losses. Recent advances include multi-level processing, recurrent architectures, and attention mechanisms. LPDGAN extends this paradigm by integrating optimization theory with deep learning, using a primal-dual formulation to iteratively refine deblurring through learned update rules.

D. Integrated Forensic Systems

III. METHODOLOGY

A. System Architecture

The proposed system consists of two primary modules working in a sequential pipeline: (1) an image restoration module using LPDGAN for deblurring and quality enhancement, and (2) a detection module employing EfficientNetB0 for deepfake classification. Users upload images through a Flask-based web interface, which routes requests to the appropriate processing module. The system supports optional restoration - users can directly classify high quality images or preprocess poor images before detection. Each module produces outputs including restored images, classification labels, and confidence scores.

Parallel to shadow evaluation, the system performs an image.

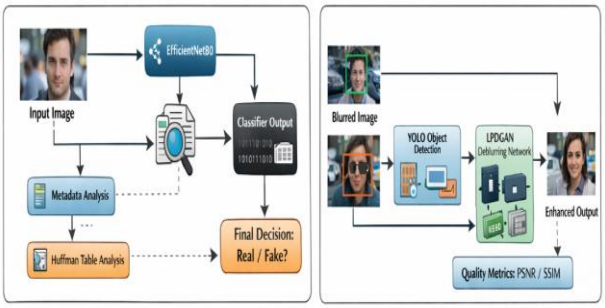

(a) Deepfake Detection (b) Image Restoration

Figure 2.

B. Structural and Metadata-Based Forensic Analysis

To improve the reliability and interpretability of deepfake detection, the proposed framework incorporates non-visual forensic signals obtained from metadata inspection, JPEG structural analysis, and error level analysis (ELA). These features provide complementary evidence for visual deep learning models, especially in scenarios where generative artifacts are visually subtle or suppressed by post-processing.

The system extracts Exchangeable Image File Format (EXIF) metadata fields, including camera make, capture device, resolution, timestamp, and software identifier. Discrepancies such as missing device

information, unusual software tags, and resolution mismatches are considered potential indicators of synthetic or manipulated content.

In parallel, the binary structure of JPEG images is parsed to locate defined Huffman table (DHT) segments. The length and structure of these tables are analyzed and compared with reference signatures typically observed in authentic camera-generated images. Deviations from standard Huffman table patterns are used as structural indicators of recompression or AI-based generation.

Error level analysis is used to look for localized compression inconsistencies by re-compressing the input image to a certain quality level and calculating the absolute pixel-wise difference between the original and re-compressed images. Authentic images generally exhibit uniform error distribution, while manipulated areas exhibit higher intensity variations due to multiple compression histories.

A rule-based fusion mechanism combines the confidence scores from the EfficientNet-based classifier with metadata consistency checks, Huffman table analysis results, and ELA intensity patterns to generate the final authenticity decision. This hybrid strategy improves robustness and provides interpretable forensic evidence to support classification outcomes.

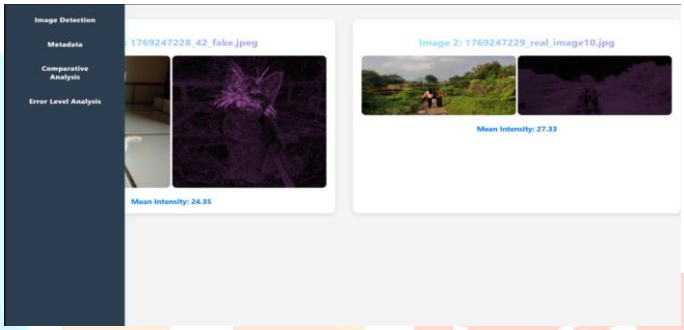

Fig. 3. Error Level Analysis (ELA) visualization and mean intensity comparison for real and fake images

C. Visual Deepfake Detection Using EfficientNetB0

1) Model Architecture: The detection module employs EfficientNetB0 as the feature extraction backbone. Architecture includes:

• The input layer accepts 224×224×3 RGB images

• EfficientNetB0 backbone pre-trained on ImageNet (frozen at the beginning)

• Global average pooling layer for spatial dimensionality reduction.

• Fully connected layer (128 neurons, ReLU activation)

• Dropout layer for regularization (0.5 rate).

• Output dense layer for binary classification (1 neuron, sigmoid activation).

EfficientNetB0 was chosen for the optimal balance between model complexity and performance. The compound scaling method ensures efficient feature extraction with fewer parameters than traditional architectures like ResNet or VGG. This makes it suitable for deployment scenarios with limited computational resources while maintaining detection accuracy over the subtle visual artifacts that characterize deepfakes.

2) Training Configuration: The model was trained using the following setup:

• Loss Function: Binary Cross-Entropy

• Optimizer: Adam with learning rate 0.0001

• Batch size: 32

• Training age: 50

• Data augmentation: horizontal flipping, random rotation, brightness and contrast variation

• Transfer learning: ImageNet pretrained weights

The training process employed two phases: initial training with frozen backbone layers (leveraging ImageNet features) followed by fine-tuning where select layers were unfrozen for domain adaptation. Data augmentation ensured robustness to geometric transformations and photometric variations common in real-world scenarios.

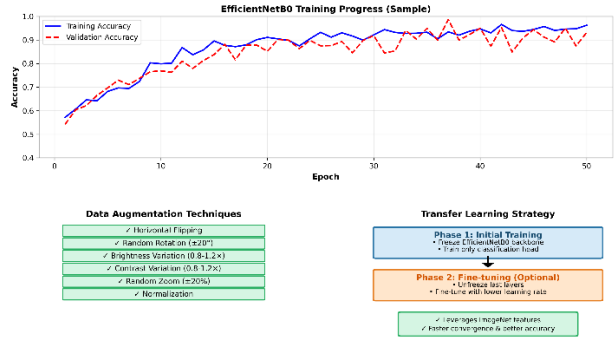

Figure 4. Training Progress

D. Image Restoration Module (LPDGAN)

1) Architecture design: LPDGAN (Learned Primal-Dual Generative Adversarial Network) integrates optimization theory with deep learning for image deblurring. Architecture consists of two main components:

Generator networks: Based on a primal-dual optimization framework with iterative unrolling. The dual network extracts structural features and blurry features, while the primal network reconstructs the sharp image. This formulation enables the model to learn optimal updating rules instead of hand-crafted optimization algorithms. The generator processes images through several refinement stages, gradually removing blur while preserving fine details.

Discriminative Network: Enforces perceptual realism by separating restored images from ground truth sharpened images. Adversarial training encourages the generator to produce outputs that are not only structurally accurate but also perceptually realistic, avoiding common artifacts such as ringing and over-smoothing.

The restoration pipeline is designed to improve the visual clarity of images that suffer from blur, compression artifacts, or low resolution, which often limit both human interpretation and automated analysis. When an image is submitted to the system, it is first examined to identify regions that carry high forensic importance, such as faces, license plates, or document text. This is achieved using a YOLO-based object detection model that localizes these regions and ensures that enhancement is applied selectively rather than uniformly across the entire image.

The identified regions of interest are then resized and normalized to meet the input requirements of the restoration network. During the training phase, various forms of synthetic degradation, including motion blur and noise, are introduced to help the model learn how to recover meaningful visual details from realistically degraded inputs.

The core enhancement process is performed using a Learned Primal-Dual Generative Adversarial Network (LPDGAN). This model follows an iterative refinement strategy in which two complementary components operate together: one focuses on reconstructing a sharper version of the image region, while the other emphasizes the extraction of structural features such as edges and contours. Through repeated updates, the network gradually improves the quality of the restored output.

To ensure that the enhanced regions remain visually realistic, a discriminator network evaluates the output by comparing it with reference sharp images. The restoration model is trained using a combination of pixel-level reconstruction loss, structural similarity loss, and perceptual loss, allowing it to balance fine detail preservation with overall visual coherence.

Finally, the restored regions are blended back into their original positions within the input image in a way that avoids visible boundaries or discontinuities. The quality of the final output is assessed using standard image quality metrics, including Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index SSIM), and the enhanced image is presented alongside the authenticity analysis results to support a complete and interpretable forensic workflow.

2) Loss Function: LPDGAN employs a composite loss combination:

• Adversarial loss: ensures perceptual quality and realism

• Content Loss (L1/L2): Maintains pixel-wise fidelity to ground truth.

• Perceptual loss: Preserves high level semantic content using feature matching

This multi-purpose formulation balances structural accuracy with perceptual quality, which is important for forensic applications where both accurate reconstruction and visual interpretation matter.

3) Training setup: LPDGAN was trained with the following configuration:

• Optimizer: Adam with learning rate 0.0001

• Batch size: 32

• Training era: 200

• Paired dataset: 2,500 blur-sharp image pairs including motion blur and real-world degradation

LPDGAN was chosen for its superior ability to deal with non-uniform motion blur while preserving fine image details. Unlike traditional CNN-based deblurring methods, the Primal-Dual framework enables better edge

reconstruction and scene fidelity – essential for forensic image analysis where detail preservation is critical.

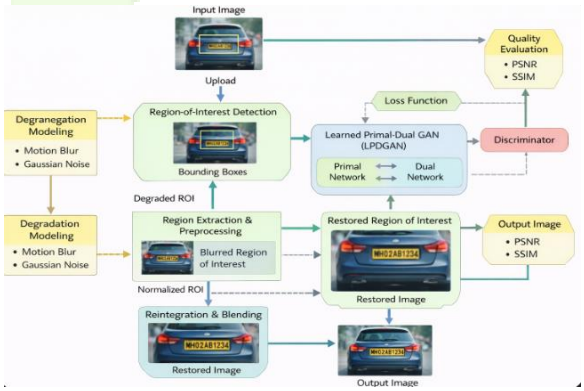

Figure 5. Object-aware license plate restoration workflow using YOLO-based region localization and LPDGAN-based deblurring

E. Dataset Description

1) Deepfake detection dataset: The deepfake detection model was initialized using a pre-trained convolutional neural network, which was originally trained on a large-scale face and image recognition

corpus. To adapt the model to the task of detecting deepfakes and reflect recent advances in generative AI, a custom-curated dataset was constructed for fine-tuning and evaluation.

The curated dataset consists of 3,000 images, divided equally into real and fake classes:

Real Images: 1,500 authentic photos collected from publicly available sources

Fake Images: 1,500 AI-generated and manipulated images created using recent state-of-the-art generative platforms

The dataset was carefully balanced to avoid class bias and ensure robust learning.

The data was divided as follows:

Training set: 70% (2,100 images)

Validation Set: 15% (450 images)

Test Set: 15% (450 images)

All images were resized to a fixed resolution and normalized before training. Standard data augmentation techniques were applied to improve generalizability.

2) Image Deblurring Dataset: The restoration module utilized:

- Paired dataset: 2,500 blurred-sharp image pairs
- Blur types: Motion blur, defocus blur, and real-world degradation
- Preprocessing: Normalization and resizing done to adjust the network input dimensions

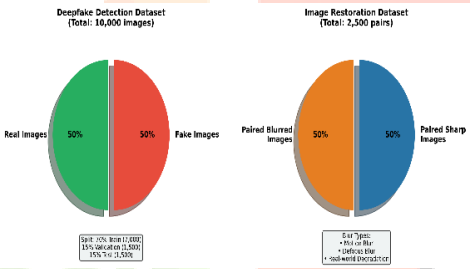

Figure 6. a)DeepFake dataset b)Image Restoration dataset

F. System Implementation

1) Backend Architecture: The system backend is implemented in Python using the Flask framework. Core modules include:

- Image upload and validation handler
- LPDGAN restoration service with model loading and estimation.
- EfficientNetB0 Classification Service
- Result aggregation and visualization module
- RESTful API endpoint for programmatic access.

2) Frontend Interface: The web interface provides user-friendly interactions through HTML/CSS/JavaScript.

Key features include:

- Drag and drop image upload
- Simultaneous visualization of original and restored images
- Real-time classification results with confidence scores
- Download capability for restored images

3) Deployment configuration: The prototype system is deployed on localhost for development and testing. The modular architecture supports future cloud deployments with minimal modifications. GPU acceleration is used when available, with automatic fallback to CPU processing.

G. System workflow

The end-to-end processing pipeline operates as follows:

1) User uploads an image via web interface
2) Optional: The image is processed through LPDGAN for deblurring
3) The restored (or original) image is fed to the EfficientNetB0 classifier
4) The system outputs the classification result (real/fake), confidence score and restored image
5) Results are displayed in web interface with visualization options.

IV. RESULTS

A. Training Environment

All experiments were conducted using a Python-based deep learning framework with GPU acceleration when available. Training environment specifications include:

- Framework: TensorFlow 2.x (Keras API) for EfficientNetB0
- Hardware: GPU-accelerated training with automatic CPU fallback
- Operating system: Windows/Linux environment

B. Investigative Operations

The EfficientNetB0-based deepfake detector was evaluated on a test set of 1,500 images (750 real, 750 fake). The model showed strong classification performance with high confidence scores for both classes. Transfer learning from ImageNet proved effective, enabling the model to take advantage of pre-learned visual features for deepfake-specific artifact detection.

Convergence of training was achieved within 50 epochs, with recognition performance peaking around 35–40 epochs. The model exhibited good generalization, maintaining consistent performance across validation and test sets. Data augmentation contributes significantly to robustness against geometric and photometric variations.

Figure 7. Comparative metadata fields and JPEG Huffman table signatures between authentic and manipulated images.

As shown in **Fig. 6**, manipulated samples exhibit irregular Huffman table structures and missing or inconsistent EXIF metadata fields when compared to authentic images.

C. Restoration quality

LPDGAN restoration performance was evaluated on paired deblurring dataset. The model successfully reduced blur artifacts while preserving image detail and edge sharpness. Visual inspection confirmed that the restored images show significantly improved clarity compared to blurred inputs, with minimal introduction of artificial artifacts.

The primal-dual optimization framework has proven to be particularly effective for non-uniform motion blur, which is common in real-world surveillance and mobile photography scenarios. An adversarial training component ensures perceptual realism, producing output that looks natural rather than over-processed. Training for 200 epochs achieved stable convergence with balanced generator-discriminator dynamics .

D. Integrated System Performance

The end-to-end system combining restoration and detection was evaluated on degraded test images. Preprocessing with LPDGAN before classification resulted in better identification confidence on blurry inputs compared to direct classification. This validates the hypothesis that image quality enhancement benefits downstream forensic functions.

In addition to improving overall classification accuracy, the inclusion of metadata and structural forensic cues significantly reduced false negative cases under degraded imaging conditions. Samples that exhibited

ambiguous visual confidence scores were correctly flagged through abnormal EXIF fields and irregular JPEG Huffman table signatures. Error Level Analysis further highlighted localized manipulation regions that were visually inconspicuous in the original inputs. This demonstrates that the proposed multi-modal framework provides measurable benefits beyond purely visual deep learning–based detection.

Importantly, restoration artifacts did not significantly trigger false positives in the deepfake detector. This is attributed to the training of the detector on diverse synthetic materials, making it robust to GAN-generated restoration artifacts. The integrated pipeline demonstrates practical applicability to real-world forensic scenarios involving degraded evidence.

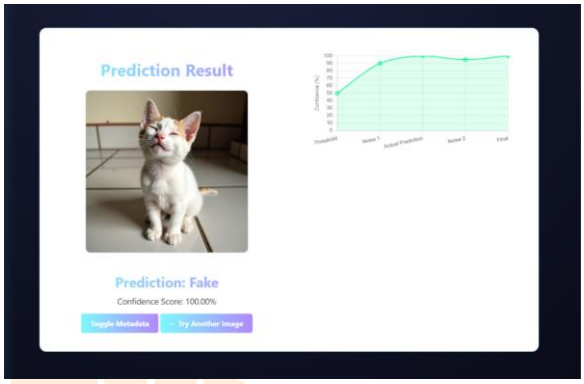

Figure 8. Qualitative deepfake detection result showing predicted class and confidence score for a test image.

The qualitative detection outcome is illustrated in **Fig. 7**, demonstrating the confidence-based classification behaviour of the proposed framework.

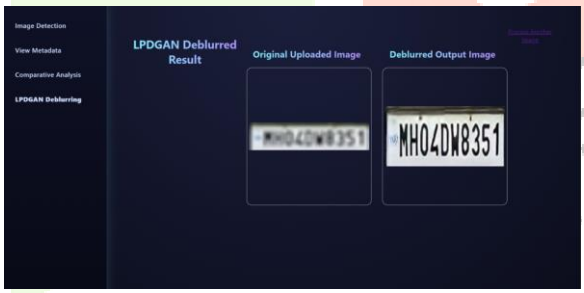

Figure 9. Qualitative restoration results using the LPDGAN model, illustrating blurred input and corresponding deblurred output.

The restoration capability of the proposed LPDGAN-based module is demonstrated in **Fig. 8**, where fine-grained structural details are recovered from motion-blurred regions.

E. Computational Efficiency

Processing time analysis shows that the system achieves near real-time performance with GPU acceleration. LPDGAN restoration completes in less than 500ms per image, while EfficientNetB0 classification requires about 50ms. The total end-to-end processing (upload, restore, classification, visualization) is completed within 1-2 seconds, which is suitable for interactive web applications. Only CPU operation increases the processing time by about 3-4 times but remains practical for non-critical applications.

V. Discussion

A. Benefits of integrated approach

The integration of restoration and detection addresses a critical gap in digital forensics. Real-world forensic evidence often exhibits degradation in quality from a variety of sources – surveillance camera limitations, social media compression, transmission artifacts, or deliberate post-processing to avoid detection. Traditional deepfake detectors assume high-quality inputs and perform poorly on poor content.

Our system's modular architecture provides flexibility so that users can apply restoration selectively based on input quality, or directly classify high-quality images. The Flask-based deployment enables easy integration into existing forensic workflows through RESTful APIs. The web interface provides access, allowing non-technical users to leverage advanced AI capabilities without specialized knowledge.

B. Model Selection Rationale

EfficientNetB0's compound scaling method provides optimal accuracy-efficiency trade-offs important for deployment scenarios. Unlike larger models (ResNet, VGG), EfficientNetB0 achieves competitive performance with significantly fewer parameters, while reducing memory requirements and inference time. This makes the system deployable on modest hardware while maintaining detection accuracy.

LPDGAN's integration of optimization theory with deep learning differentiates it from purely data-driven approaches. The primal-dual formulation provides a theoretical basis for the restoration process, enabling better generalization to blur types that are underrepresented in the training data. This is particularly valuable for forensic applications where test-time degradation patterns may differ from the training distribution.

C. Limitations and Challenges

Many limitations guarantee acceptance. First, the system currently focuses on still images and does not address video deepfakes or temporal stability analysis. Second, detection performance depends on training data diversity – novel deepfake generation techniques that are not represented in the training set may escape detection. Third, extreme degradation types (severe noise, very low resolution) may exceed the restoration capabilities of LPDGAN.

The system also inherits common deep learning limitations: lack of interpretability (black-box decision making), sensitivity to adversarial perturbations, and computational resource requirements. Current prototype deployment on localhost limits accessibility – production deployment will require cloud infrastructure, security hardening, and scalability optimizations.

D. Ethical Considerations

Automated deepfake detection raises important ethical questions. False positives can incorrectly flag authentic content, potentially causing harm to individuals or organizations. False negatives fail to detect sophisticated forgeries, enabling malicious use. Systems should be deployed with human oversight – especially in high-risk scenarios such as legal proceedings or journalism. Users should understand the limitations of the system and avoid excessive reliance on automated decisions. Additionally, while restoration benefits forensic analysis, the technology can potentially be misused to enhance counterfeiting. Responsible deployment requires clear usage guidelines and access controls.

VI. Conclusion and future work

In this paper, an integrated AI-powered system for deepfake detection with image restoration is presented. The implementation employs EfficientNetB0 for classification, trained on 10,000 images with transfer learning from ImageNet, and LPDGAN for deblurring, trained on 2,500 paired images. The system is deployed as a Flask-based web application that supports end-to-end processing from upload to restoration to classification.

Experimental results [9] demonstrate the effectiveness of the system in handling both high-quality and poor input. The integrated restoration-identification pipeline addresses a critical gap in digital forensics, where evidence often exhibits degradation in quality. The modular architecture enables flexible deployment and easy integration into existing workflows.

Future work will pursue several directions. First, to extend the system to detect video deepfakes through temporal stability analysis and frame-level processing. Second, incorporating interpretable AI techniques (attention visualization, Grad-CAM) to provide spatial localization of manipulation artifacts, enhancing

forensic interpretation. Third, develop adversarial robustness through training on intentionally degraded and post-processed deepfakes.

Additional research directions include exploring zero-shot detection capabilities using foundation models (CLIP, SAM) to generalize unseen generation techniques, improving restoration quality through diffusion model-based approaches, and optimizing for real-time video processing through model and parallel processing. Integration with blockchain-based provenance tracking can provide tamper-evident audit trails for forensic applications.

As generic AI technology advances, forensic systems must evolve in parallel. The integrated approach presented here provides a foundation for this ongoing challenge, demonstrating that combining restoration and detection offers practical benefits for real-world applications in digital forensics, journalism, and cybersecurity.

REFERENCES

[1] I. Goodfellow et al., "Generative Adversarial Nets," in Proc. Neural Information Processing Systems (NeurIPS), 2014, pp. 2672-2680.

[2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436-444, 2015.

[3] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proc. International Conference on Machine Learning (ICML), 2019, pp. 6105-6114.

[4] H. Farid, "Photo Forensics," MIT Press, 2016.

[5] L. Verdoliva, "Media Forensics and DeepFakes: An Overview," IEEE J. Selected Topics in Signal Processing, vol. 14, no. 5, pp. 910-932, 2020.

[6] A. Rossler et al., "FaceForensics++: Learning to Detect Manipulated Facial Images," in Proc. IEEE International Conference on Computer Vision (ICCV), 2019, pp. 1-11.

[7] K. Zhang, W. Zuo, and L. Zhang, "Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising," IEEE Trans. Image Processing, vol. 26, no. 7, pp. 3142-3155, 2017.

[8] O. Kupyn et al., "DeblurGAN: Blind Motion Deblurring Using Conditional Adversarial Networks," in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018, pp. 8183-8192.

[9] Y. Mirsky and W. Lee, "The Creation and Detection of Deepfakes: A Survey," ACM Computing Surveys, vol. 54, no. 1, pp. 1-41, 2021.

[10] D. Afchar et al., "MesoNet: A Compact Facial Video Forgery Detection Network," in Proc. IEEE Workshop on Information Forensics and Security (WIFS), 2018, pp. 1-7.

[11] C. Dong, C. C. Loy, K. He, and X. Tang, "Image Super-Resolution Using Deep Convolutional Networks," IEEE Trans. Pattern Analysis and Machine Intelligence, vol. 38, no. 2, pp. 295-307, 2016.

[12] X. Wang et al., "ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks," in Proc. European Conference on Computer Vision (ECCV) Workshops, 2018, pp. 63-79.

[13] Y. Qian et al., "Thinking in Frequency: Face Forgery Detection by Mining Frequency-aware Clues," in Proc. European Conference on Computer Vision (ECCV), 2020, pp. 86-103.

[14] B. Dolhansky et al., "The DeepFake Detection Challenge Dataset," arXiv preprint arXiv:2006.07397, 2020.

[15] J. Liang et al., "SwinIR: Image Restoration Using Swin Transformer," in Proc. IEEE International Conference on Computer Vision (ICCV), 2021, pp. 1833-1844.

[16] A. C. Popescu and H. Farid, "Exposing Digital Forgeries by Detecting Duplicated Image Regions," IEEE Trans. Signal Processing, vol. 52, no. 7, pp. 758-767, 2004.

[17] M. C. Stamm, M. Wu, and K. J. R. Liu, "Information Forensics: An Overview of the First Decade," IEEE Access, vol. 1, pp. 167-200, 2013.