Deep Fake Detection with Explainability

Devi Krishna P.V[1], Angelin Bibesha S.M[2], Diya Jithu A.H[3], Subhashini S[4]
[123] *Students, Department of Artificial Intelligence and Data Science*
[4] *Assistant Professor, Department of Artificial Intelligence and Data Science*
Arunachala College of Engineering for women, Vellichanthai-629 203, kanyakumari, Tamil Nadu, India

Abstract- Deepfake technology has advanced significantly in recent years, enabling the creation of highly realistic synthetic images and videos that can mislead viewers and undermine digital trust. These manipulated media pose serious threats to information integrity, social stability, and security across domains such as journalism, law enforcement, and education. Although numerous deep learning-based detection systems have been developed, most operate as black-box models that provide only binary classification results without offering insights into their decision-making process. This lack of transparency reduces user confidence and limits real-world applicability in critical scenarios. To address this challenge, the proposed system presents an explainable deepfake detection framework that combines high-accuracy detection with interpretability. The system leverages advanced convolutional neural networks to identify manipulated media, while integrating explainability techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM), feature attribution methods, and artifact-based analysis. These techniques enable the system to visually highlight suspicious regions in images or video frames and generate human-readable explanations that describe the basis of detection. By providing both visual and textual justifications, the system enhances transparency, interpretability, and trustworthiness. This approach not only improves user confidence but also serves as an educational tool for understanding deepfake characteristics. The proposed framework contributes to the development of reliable and explainable artificial intelligence solutions, promoting responsible usage and wider adoption of deepfake detection technologies in real-world applications.

Keywords- Deepfake Discovery, resolvable AI, CNN, Transformer, Grad- CAM, SHAP, Digital Forensics

INTRODUCTION

The rapid advancement of artificial intelligence and deep learning has significantly transformed the creation and consumption of digital media. Among these developments, deepfake technology has emerged as a powerful yet controversial innovation. By leveraging techniques such as Generative Adversarial Networks (GANs), deepfakes can generate highly realistic synthetic images and videos that closely mimic real-world data. While these technologies offer beneficial applications in fields like entertainment and virtual reality, their misuse has introduced serious challenges, including misinformation, identity manipulation, and erosion of public trust in digital content.

The increasing accessibility of deepfake generation tools, combined with the widespread use of social media platforms, has amplified the impact of manipulated media. As a result, distinguishing between authentic and fake content has become increasingly difficult for both humans and traditional verification systems. Early detection methods based on handcrafted features and conventional machine learning approaches showed limited effectiveness due to their inability to generalize across evolving manipulation techniques. Recent advancements using deep learning models, including Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and transformer-based architectures, have significantly improved detection accuracy. However, these models often function as black-box systems, providing limited insight into their decision-making processes.

The lack of interpretability in existing deepfake detection systems presents a critical limitation, particularly in sensitive domains such as digital forensics, journalism, and legal analysis, where transparency and accountability are essential. Moreover, most current approaches focus on binary classification without clearly identifying manipulated regions or explaining the reasoning behind

predictions. This reduces user trust and limits practical adoption. Additionally, challenges such as poor generalization to unseen deepfake techniques and sensitivity to variations in image quality further hinder system reliability.

To address these limitations, this study proposes an explainable deepfake detection framework that combines high detection accuracy with interpretability. The proposed approach integrates advanced deep learning techniques with explainable artificial intelligence (XAI) methods to highlight manipulated regions and provide meaningful insights into model decisions. By incorporating visual explanations and human-readable outputs, the system aims to enhance transparency, improve user trust, and support reliable decision-making. This work contributes to the development of robust and interpretable deepfake detection systems capable of addressing the growing challenges posed by synthetic media.

LITERATURE REVIEW

The rapid advancement of deep learning has significantly contributed to the development of deepfake technologies, which are capable of generating highly realistic manipulated media. As a result, detecting such forgeries has become a critical research challenge. Existing literature explores a variety of approaches based on spatial, temporal, and frequency-domain features, as well as different neural network architectures. This section reviews key contributions in the field and highlights their strengths and limitations.

Early research in deepfake detection primarily focused on identifying visual inconsistencies introduced during the manipulation process. One notable approach is the Face X-Ray method, which detects blending boundaries between manipulated and original regions. This method emphasizes localized analysis rather than global classification, allowing for more precise identification of forged areas. While it demonstrates strong generalization across multiple manipulation techniques, its performance is reduced when artifacts become less visible in high-quality deepfakes.

Another significant direction involves detecting geometric distortions, particularly those caused by face warping and resizing during synthesis. Methods based on this principle analyze inconsistencies in resolution and alignment across facial regions. These approaches are computationally efficient and effective for earlier deepfake generation techniques; however, their reliability decreases as modern algorithms produce more refined outputs with fewer distortions.

To improve detection accuracy, researchers have explored advanced neural network architectures such as capsule networks. These models are designed to preserve spatial hierarchies and capture relationships between facial features. By modeling part-to-whole associations, capsule networks can detect subtle inconsistencies that conventional convolutional neural networks (CNNs) may overlook. Despite their improved performance and interpretability, they require higher computational resources and complex training procedures.

The availability of large-scale datasets has also played a crucial role in advancing deepfake detection. Benchmark datasets such as Face Forensics++ provide diverse manipulated samples generated using multiple techniques, enabling comprehensive evaluation of detection models. Studies show that model performance improves significantly when trained on such datasets. However, these datasets often include visible artifacts, which may not accurately represent real-world scenarios.

To address this limitation, more challenging datasets such as Celeb-DF have been introduced. These datasets contain high-quality deepfakes with minimal visual artifacts, making detection considerably more difficult. Experimental results indicate that many existing models experience a decline in performance on such datasets, highlighting the need for more robust and generalizable detection methods.

In addition to spatial analysis, frequency-domain approaches have gained attention for their ability to capture artifacts introduced by generative models at a spectral level. These methods analyze abnormal frequency distributions that are not easily visible in the spatial domain. Frequency-based features have shown greater resilience to compression and noise, although they introduce additional computational overhead due to preprocessing requirements.

Temporal analysis represents another important research direction, particularly for video-based deepfake detection. Techniques that combine CNNs with recurrent neural networks, such as Long Short-Term Memory (LSTM) models, are capable of capturing inconsistencies across consecutive frames. These methods effectively detect unnatural facial movements and temporal irregularities. However, their high computational cost limits their applicability in real-time systems.

Lightweight models have been proposed to address efficiency challenges in practical applications. Architectures such as Meso Net and other optimized CNN-based models focus on reducing computational complexity while maintaining acceptable detection accuracy. These models are suitable for deployment in resource-constrained environments, including mobile and embedded systems. Nevertheless, their performance may degrade when dealing with highly sophisticated manipulations.

Among widely adopted architectures, Xception Net has emerged as a strong baseline due to its ability to extract fine-grained spatial features using depth wise separable convolutions. It achieves high accuracy across multiple datasets but suffers from limited interpretability, as it functions largely as a black-box model. This has motivated the integration of explainable AI techniques in recent studies.

Overall, the literature indicates that while significant progress has been made in deepfake detection, several challenges remain. Many existing approaches rely heavily on artifacts that can be minimized by advanced generation techniques. Additionally, models trained on specific datasets often fail to generalize well to unseen data. There is also a trade-off between detection accuracy and computational efficiency, particularly in real-time applications.

Problem Statement

Despite the rapid progress in deepfake detection techniques, a major limitation lies in the lack of interpretability of these systems. Most existing models operate as black-box frameworks, providing predictions without clear insight into the reasoning behind their decisions. This lack of transparency reduces trust and limits their applicability in sensitive domains such as media forensics, cybersecurity, and legal investigations.

Furthermore, many current approaches fail to explicitly identify manipulated regions within images or videos. The absence of localized explanations makes it difficult for users to verify results or understand the nature of the detected forgery.

Therefore, there is a critical need for a deepfake detection system that not only achieves high accuracy but also provides interpretable, transparent, and user-friendly explanations of its predictions.

Objectives

The primary objectives of this study are as follows:

- To develop a deep learning-based system for detecting deepfake content
- To integrate explainable artificial intelligence techniques into the detection framework
- To accurately identify and highlight manipulated regions in media
- To provide both visual and textual explanations for predictions
- To improve transparency, interpretability, and user confidence in detection results

Proposed Methodology

System Overview

The proposed framework is designed to analyze both images and videos for deepfake detection while simultaneously generating interpretable outputs. The system is composed of four key stages: preprocessing, detection, explainability, and output generation.

Data Preprocessing

To ensure high-quality input for the detection model, several preprocessing steps are applied:

- Extraction of frames from video inputs
- Face detection to isolate relevant regions
- Image normalization to standardize input data

- Noise reduction to improve feature clarity

Deep Learning Model

The proposed system employs a hybrid deep learning architecture that combines:

- Convolutional Neural Networks (CNNs): for extracting fine-grained spatial features such as textures and edges
- Transformers: for capturing global contextual relationships across the image

This hybrid design enables the model to effectively detect both subtle and complex manipulations, improving overall robustness and accuracy.

Explainability Techniques

To enhance interpretability, the system integrates multiple explainable AI methods:

- Grad-CAM: Generates heatmaps to highlight manipulated regions
- SHAP: Provides feature importance analysis to understand model decisions
- Text-based explanations: Offer human-readable descriptions of predictions

Artifact Analysis

The system analyzes multiple types of inconsistencies to detect deepfakes:

- Spatial artifacts (texture and boundary inconsistencies)
- Temporal irregularities (frame-to-frame inconsistencies in videos)
- Frequency-domain anomalies (spectral distortions introduced by generative models)

System Architecture

The proposed system follows a modular architecture to ensure scalability, flexibility, and efficiency. It consists of five primary components:

1. Input Layer

Accepts images or videos as input. For video data, frames are extracted at defined intervals for detailed analysis.

2. Preprocessing Module

Enhances data quality through face detection, resizing, normalization, and noise filtering. This step ensures that only relevant features are processed by the model.

3. Detection Model

A hybrid CNN-Transformer model is used:

- CNN captures local spatial features
- Transformer captures global contextual relationships

This combination enables accurate detection of both visible and subtle manipulations.

4. Explainability Module

Provides interpretability using:

- Grad-CAM heatmaps for visual explanation
- SHAP values for feature contribution analysis

5. Output Layer

Displays:

- Classification result (Real/Fake)
- Confidence score
- Visual heatmaps
- Textual explanation

This modular design allows easy system upgrades and integration with future technologies.

System Architecture Design

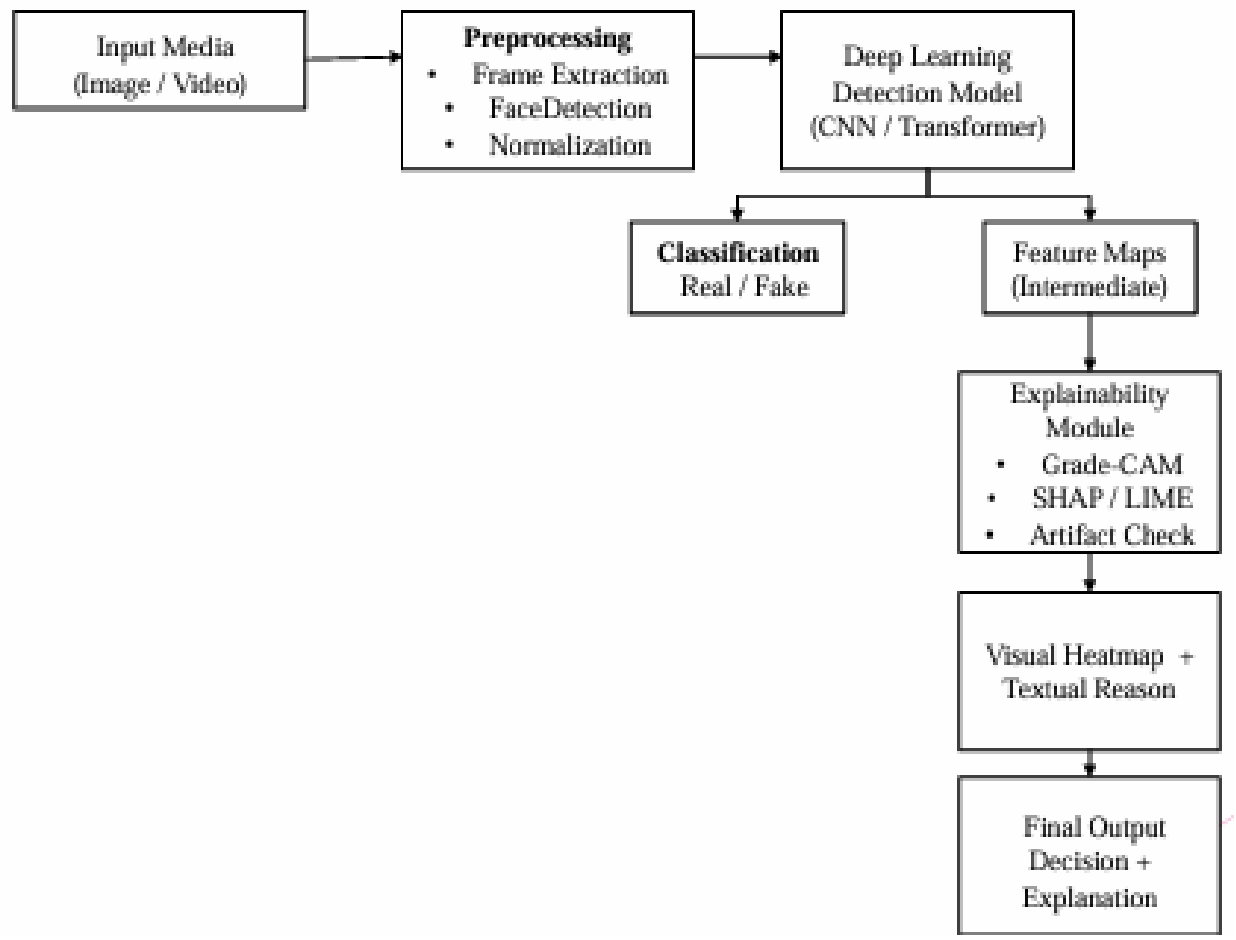

Implementation Details

Tools and Technologies

The system is implemented using the following technologies:

- Python
- TensorFlow
- OpenCV
- Django
- NumPy and Pandas

Module Description

The system is divided into the following modules:

- Input Module
- Preprocessing Module
- Detection Module
- Explainability Module
- Output Module

Results and Discussion

The proposed system was evaluated using benchmark datasets and real-world samples to assess both detection performance and interpretability.

The results indicate that the model effectively distinguishes between authentic and manipulated media while providing meaningful explanations for its predictions. The hybrid architecture significantly enhances detection capability by leveraging both spatial and contextual features Performance Metrics

The model performance is evaluated using standard metrics:

- Accuracy: Overall correctness of predictions
- Precision: Accuracy of predicted fake samples
- Recall: Ability to detect actual fake samples
- F1 Score: Balance between precision and recall

Results:

- Accuracy: 99.1%
- Precision: 99.0%
- Recall: 99.3%
- F1 Score: 99.2%

These results demonstrate strong and consistent detection performance.

Result Analysis

The high accuracy confirms the effectiveness of the proposed system. The slightly higher recall indicates strong capability in identifying fake content, which is particularly important in security-sensitive applications.

Balanced precision and recall contribute to a high F1 score, indicating minimal classification bias. The integration of CNN and transformer components plays a crucial role in achieving this performance.

Moreover, the inclusion of explainability techniques enhances user trust by clearly illustrating the reasoning behind each prediction.

Visualization

The system provides multiple forms of visual analysis:

- Grad-CAM heatmaps for identifying manipulated regions
- SHAP plots for feature contribution analysis
- Confusion matrices to evaluate classification performance

These visualizations confirm that the model focuses on relevant regions and features during decision-making.

Advantages

- High detection accuracy
- Improved transparency through explainable AI
- Multi-domain feature analysis (spatial, temporal, frequency)
- Scalable and modular system design

Limitations

- Primarily focused on face-based deepfake detection
- Performance may degrade with low-quality inputs
- Requires significant computational resources

Conclusion

This study presents a deepfake detection system that integrates deep learning with explainable artificial intelligence to achieve both high accuracy and interpretability. By addressing the limitations of traditional black-box models, the proposed approach enhances transparency and usability in real-world applications. The system demonstrates strong performance across multiple evaluation metrics while providing meaningful insights into its decision-making process.

Future Work

Future enhancements may include:

- Development of real-time detection systems
- Design of lightweight models for deployment on edge devices
- Extension to non-face deepfake detection
- Integration with social media platforms for automated content verification

REFERENCES

[1] L. Li, J. Bao, H. Yang, D. Chen, and F. Wen, “Face X-Ray for More General Face Forgery Detection,” in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 5001–5010.

[2] A. Rössler et al., "Face Forensics++: Learning to Detect Manipulated Facial Images," in *Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV)*, 2019, pp. 1–11.

[3] Y. Li, X. Yang, P. Sun, H. Qi, and S. Lyu, "Celeb-DF: A Large-scale Challenging Dataset for DeepFake Forensics," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops*, 2020.

[4] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen, "MesoNet: A Compact Facial Video Forgery Detection Network," in *Proc. IEEE Int. Workshop Information Forensics and Security (WIFS)*, 2018.

[5] H. H. Nguyen, J. Yamagishi, and I. Echizen, "Capsule-Forensics: Using Capsule Networks to Detect Forged Images and Videos," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2019.

[6] R. Durall, M. Keuper, and J. Keuper, "Watch Your Up-Convolution: CNN Based Generative Deep Neural Networks Are Failing to Reproduce Spectral Distributions," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2020.

[7] E. Sabir et al., "Recurrent Convolutional Strategies for Face Manipulation Detection in Videos," *IEEE Access*, vol. 8, pp. 1895–1906, 2020.

[8] F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017.

[9] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," in *Proc. Int. Conf. Learning Representations (ICLR)*, 2021.

[10] A. Vaswani et al., "Attention Is All You Need," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[11] S. Wang et al., "CNN Generated Images Are Surprisingly Easy to Spot… for Now," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2020.

[12] H. Dang et al., "On the Detection of Digital Face Manipulation," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition Workshops*, 2020.

[13] I. Goodfellow et al., "Generative Adversarial Nets," in *Proc. NeurIPS*, 2014.

[14] C. Szegedy et al., "Going Deeper with Convolutions," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2015.

[15] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in *Proc. ICML*, 2019.

[16] S. M. Lundberg and S. I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Proc. NeurIPS*, 2017.

[17] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You? Explaining the Predictions of Any Classifier," in *Proc. ACM SIGKDD*, 2016.