Hybrid CNN–Transformer Framework for Robust Deepfake Detection

Shreya Gupta[1], Krish Lal Srivastava[2], Diksha Asnora[3], Harshit Mehra[4], Mr. Abhishek Chaudhary[5]
[1,2,3,4,5]*Department of Computer Science and Engineering, MGM's College of Engineering and Technology, Noida, Uttar Pradesh, India*

***Abstract*— The rapid advancement of generative artificial intelligence has intensified the challenge of distinguishing authentic media from synthetic deepfakes. Building upon prior CNN-based frameworks, this continuation research explores advanced detection strategies that address real-world deployment issues such as adversarial robustness, multimodal manipulations, and interpretability. A hybrid architecture integrating convolutional neural networks with transformer-based attention mechanisms is proposed, alongside multimodal fusion of image, audio, and metadata features. Experimental evaluation across benchmark datasets demonstrates improved generalization, resilience against compression, and enhanced transparency through explainable AI techniques. The findings highlight that scalable, interpretable, and fairness-aware detection systems are essential for practical applications in journalism, law enforcement, and social media moderation, thereby contributing to the broader goal of safeguarding digital trust in the era of synthetic media.**

***Index Terms*—Deepfake Detection, CNN, Transformer, Explainable AI, Multimodal Fusion, Adversarial Robustness**

I. INTRODUCTION

The continuous evolution of artificial intelligence has transformed the landscape of digital media, enabling the creation of synthetic content that is increasingly difficult to distinguish from authentic sources. Deepfake technology, powered by advanced generative models such as GANs and diffusion networks, has progressed beyond simple facial swaps to highly realistic manipulations that integrate audio, video, and textual elements. While the first phase of research primarily focused on convolutional neural networks (CNNs) for image-based detection, the growing sophistication of deepfakes necessitates more advanced and adaptable approaches.

This second study builds upon earlier CNN-based frameworks by addressing challenges encountered in real-world deployment, including adversarial robustness, multimodal manipulation, and interpretability. The misuse of deepfakes continues to pose risks in domains such as misinformation, identity theft, political propaganda, and digital forensics. Consequently, the development of scalable, transparent, and fairness-aware detection systems has become a critical priority.

By exploring hybrid architectures that combine CNNs with transformer-based attention mechanisms, and by incorporating multimodal fusion strategies, this research aims to extend the capabilities of deepfake detection beyond controlled datasets. The emphasis is placed on enhancing generalization, improving resilience against compression and adversarial attacks, and providing interpretable outputs that strengthen trust in automated forensic systems. This continuation study therefore contributes to bridging the gap between theoretical advancements and practical deployment in diverse, real-world environments.

II. LITERATURE REVIEW

The detection of deepfake content has become one of the most critical research areas in artificial intelligence and digital forensics. Early approaches primarily relied on handcrafted features such as texture inconsistencies, lighting variations, and compression artifacts. While these methods provided reasonable accuracy for simple manipulations, they failed to generalize against high-quality deepfakes generated by advanced neural networks.

With the rise of deep learning, convolutional neural networks (CNNs) emerged as a dominant technique for deepfake detection. Afchar et al. introduced

MesoNet, a lightweight CNN architecture that analyzed mesoscopic features of facial images. Although effective for certain manipulations, MesoNet struggled with large-scale datasets and high-resolution synthetic images. Rössler et al. later developed the Face Forensics++ dataset, which became a benchmark for evaluating detection models. Using this dataset, deeper architectures such as XceptionNet achieved high accuracy by learning complex spatial features, though at the cost of significant computational resources.

Beyond CNNs, researchers explored alternative strategies to improve robustness. Capsule networks demonstrated the ability to capture spatial hierarchies, while frequency-domain analysis revealed subtle GAN-related artifacts invisible in the spatial domain. Transformer-based architectures introduced multi-head attention mechanisms, enabling models to capture long-range dependencies and outperform CNNs in high-resolution scenarios. Additionally, multimodal approaches that combined visual, audio, and textual cues showed promise in detecting manipulations across diverse media formats.

Explainability has also become a major focus in recent literature. Techniques such as Grad-CAM and SHAP provide visual explanations of model decisions, enhancing transparency and trust in forensic applications. However, challenges remain in ensuring cross-dataset generalization, resilience against adversarial attacks, and fairness across demographic groups.

Overall, the literature highlights a clear progression from handcrafted features to deep learning, and now toward hybrid and multimodal architectures. This evolution underscores the need for detection systems that balance accuracy, efficiency, and interpretability, making them suitable for deployment in real-world environments where deepfakes continue to evolve rapidly.

III. PROBLEM STATEMENT

Despite significant progress in deepfake detection research, several unresolved challenges continue to hinder the effectiveness of existing solutions in real-world environments. As generative models evolve from traditional GANs to diffusion-based architectures and multimodal synthesis, detection systems must adapt to increasingly sophisticated manipulations while maintaining efficiency and reliability.

One major limitation is cross-dataset generalization. Many detection models perform well on specific datasets but fail when exposed to unseen manipulations or compressed social media content. This reduces their applicability in diverse, real-world scenarios.

Another critical issue is robustness against adversarial attacks. Malicious actors can deliberately introduce perturbations that deceive detection systems, undermining their reliability in forensic and security applications.

Multimodal manipulation further complicates detection. Deepfakes are no longer limited to facial images; they now integrate audio, video, and textual cues, requiring detection systems to analyze multiple modalities simultaneously.

Interpretability also remains a challenge. Most deep learning models function as black boxes, providing predictions without clear justification. For forensic and investigative use, it is essential that detection systems offer transparent and explainable outputs.

Finally, fairness and bias in datasets affect model reliability across demographics. Unequal representation in training data can lead to biased predictions, raising ethical concerns in sensitive applications.

The objective of this continuation research is to design a hybrid deepfake detection framework that integrates CNNs with transformer-based architectures, incorporates multimodal fusion, and emphasizes adversarial robustness, interpretability, and fairness. This approach aims to bridge the gap between controlled experimental accuracy and practical deployment in diverse, real-world conditions.

IV. METHODOLOGY

The methodology adopted in this continuation research builds upon the CNN-based framework introduced in the first study, while extending it with hybrid architectures and multimodal strategies to address real-world challenges. The process is structured into sequential stages: dataset acquisition, preprocessing, model design, training, and evaluation.

A. Data Acquisition and Preprocessing

Benchmark datasets such as Face Forensics++, Celeb-DF, DFDC, and FakeAVCeleb were utilized to ensure

diversity in manipulations and modalities. Facial regions were extracted using automated detection and alignment, resized to a fixed resolution of 224 × 224 pixels, and normalized for numerical stability. To simulate real-world conditions, preprocessing included compression artifacts, noise injection, and multimodal alignment of audio and metadata. Data augmentation techniques such as rotation, flipping, brightness variation, and adversarial perturbations were applied to improve robustness.

B. Hybrid Model Architecture

The proposed detection system integrates convolutional neural networks (CNNs) for local feature extraction with transformer-based attention mechanisms for capturing long-range dependencies. CNN layers focus on texture and pixel-level artifacts, while transformer layers analyze global relationships across facial regions. For multimodal fusion, audio spectrograms and textual metadata are processed alongside visual features, enabling detection of cross-modal manipulations.

C. Training Strategy

The model was trained using binary cross-entropy loss with the Adam optimizer and a low learning rate to ensure stable convergence. Transfer learning was employed by initializing CNN layers with pretrained ImageNet weights, while transformer layers were fine-tuned on deepfake datasets. Adversarial training was incorporated to strengthen resilience against perturbations. The dataset was split into training, validation, and testing subsets in a 70:15:15 ratio, with early stopping applied to prevent overfitting.

D. Evaluation Metrics

Performance was assessed using accuracy, precision, recall, F1-score, and area under the ROC curve (AUC). Cross-dataset testing was conducted to evaluate generalization, while adversarial robustness was measured by introducing controlled perturbations. Interpretability was analyzed using Grad-CAM and SHAP visualizations to highlight manipulated regions influencing classification decisions.

This methodology ensures that the proposed framework not only achieves high detection accuracy but also addresses critical challenges of robustness, scalability, and transparency, making it suitable for deployment in real-world forensic and media authenticity applications.

Table I. Comparative Overview of Deepfake Detection Challenges and Proposed Solutions

Challenge	Description	Limitations in Existing Models	Proposed Solution in Continuation Framework
Cross-Dataset Generalization	Models fail when tested on unseen datasets or new manipulation techniques.	Accuracy drops significantly outside training dataset.	Hybrid CNN-Transformer architecture with transfer learning for adaptability.
Robustness to Compression	Social media platforms apply compression, obscuring forensic artifacts.	CNN-based detectors lose reliability under heavy compression.	Data augmentation with simulated compression and adversarial training.
Adversarial Attacks	Malicious perturbations deceive detection systems.	Vulnerable to small pixel-level changes.	Incorporation of adversarial defense strategies and robust augmentation.
Multimodal Manipulation	Deepfakes now integrate audio, video, and text cues.	Image-only models cannot detect multimodal manipulations.	Fusion of visual, audio, and metadata features for multimodal detection.
Interpretability	Deep learning models act as black boxes, limiting forensic trust.	Predictions lack transparency and justification.	Explainable AI techniques (Grad-CAM, SHAP) for visual and feature-level insights.
Fairness and Bias	Dataset imbalance reduces reliability across demographics.	Biased predictions in underrepresented groups.	Fairness-aware training with diverse datasets and balanced representation.

This continuation table highlights how the proposed framework extends beyond the first research paper by addressing real-world deployment challenges. It emphasizes generalization, robustness, multimodal

fusion, interpretability, and fairness, ensuring that the detection system is scalable and trustworthy in diverse environments.

V. SYSTEM DESIGN

The system comprises three modules: (1) Preprocessing and feature extraction using CNNs, (2) Transformer-based temporal modeling, and (3) Decision fusion and explainability layer. The architecture is optimized for real-time inference and interpretability. The methodology adopted in this continuation research builds upon the CNN-based framework introduced in the first study, while extending it with hybrid architectures and multimodal strategies to address real-world challenges. The process is structured into sequential stages: dataset acquisition, preprocessing, model design, training, and evaluation.

A. Data Acquisition and Preprocessing

Benchmark datasets such as FaceForensics++, Celeb-DF, DFDC, and FakeAVCeleb were utilized to ensure diversity in manipulations and modalities. Facial regions were extracted using automated detection and alignment, resized to a fixed resolution of 224 × 224 pixels, and normalized for numerical stability. To simulate real-world conditions, preprocessing included compression artifacts, noise injection, and multimodal alignment of audio and metadata. Data augmentation techniques such as rotation, flipping, brightness variation, and adversarial perturbations were applied to improve robustness.

B. Hybrid Model Architecture

The proposed detection system integrates convolutional neural networks (CNNs) for local feature extraction with transformer-based attention mechanisms for capturing long-range dependencies. CNN layers focus on texture and pixel-level artifacts, while transformer layers analyze global relationships across facial regions. For multimodal fusion, audio spectrograms and textual metadata are processed alongside visual features, enabling detection of cross-modal manipulations.

C. Training Strategy

The model was trained using binary cross-entropy loss with the Adam optimizer and a low learning rate to ensure stable convergence. Transfer learning was employed by initializing CNN layers with pretrained ImageNet weights, while transformer layers were fine-tuned on deepfake datasets. Adversarial training was incorporated to strengthen resilience against perturbations. The dataset was split into training, validation, and testing subsets in a 70:15:15 ratio, with early stopping applied to prevent overfitting.

D. Evaluation Metrics

Performance was assessed using accuracy, precision, recall, F1-score, and area under the ROC curve (AUC). Cross-dataset testing was conducted to evaluate generalization, while adversarial robustness was measured by introducing controlled perturbations. Interpretability was analyzed using Grad-CAM and SHAP visualizations to highlight manipulated regions influencing classification decisions.

This methodology ensures that the proposed framework not only achieves high detection accuracy but also addresses critical challenges of robustness, scalability, and transparency, making it suitable for deployment in real-world forensic and media authenticity applications.

Table I. Comparative Overview of Deepfake Detection Challenges and Proposed Solutions.

Challenge	Description	Limitations in Existing Models	Proposed Solution in Continuation Framework
Cross-Dataset Generalization	Models fail when tested on unseen datasets or new manipulation techniques.	Accuracy drops significantly outside training dataset.	Hybrid CNN-Transformer architecture with transfer learning for adaptability.
Robustness to Compression	Social media platforms apply compression, obscuring forensic artifacts.	CNN-based detectors lose reliability under heavy compression.	Data augmentation with simulated compression and adversarial training.
Adversarial Attacks	Malicious perturbations deceive detection systems.	Vulnerable to small pixel-level changes.	Incorporation of adversarial defense strategies and robust augmentation.

Challenge	Description	Limitations in Existing Models	Proposed Solution in Continuation Framework
Multimodal Manipulation	Deepfakes now integrate audio, video, and text cues.	Image-only models cannot detect multimodal manipulations.	Fusion of visual, audio, and metadata features for multimodal detection.
Interpretability	Deep learning models act as black boxes, limiting forensic trust.	Predictions lack transparency and justification.	Explainable AI techniques (Grad-CAM, SHAP) for visual and feature-level insights.
Fairness and Bias	Dataset imbalance reduces reliability across demographics.	Biased predictions in underrepresented groups.	Fairness-aware training with diverse datasets and balanced representation.

VI. RESULTS AND DISCUSSION

This section presents the evaluation of the proposed hybrid CNN–Transformer based deepfake detection framework. The results are analyzed using standard classification metrics and compared with existing models to highlight improvements in accuracy, robustness, and interpretability.

A. Quantitative Performance Analysis

The framework was tested on benchmark datasets including FaceForensics++, Celeb-DF, DFDC, and FakeAVCeleb. Performance metrics such as accuracy, precision, recall, and F1-score were computed.

Table II. Performance of Proposed Hybrid Model Across Datasets

Dataset	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
FaceForensics++	97.2	96.8	96.4	96.6
Celeb-DF	96.5	95.9	95.2	95.5
DFDC	93.8	92.7	92.1	92.4
FakeAVCeleb (Multimodal)	94.6	93.9	93.2	93.5

The results indicate that the hybrid model consistently outperforms traditional CNN-only architectures, particularly in cross-dataset generalization and multimodal detection.

B. Comparative Analysis

To evaluate efficiency, the proposed framework was compared with existing models such as MesoNet, XceptionNet, and Capsule Networks.

Table III. Comparative Performance of Detection Models

Model	Accuracy (%)	Inference Time (ms)	Interpretability
MesoNet	92.4	15	Limited
XceptionNet	94.6	45	Moderate
Capsule Networks	93.1	60	Limited
Proposed Hybrid CNN–Transformer	97.2	28	High (Grad-CAM, SHAP)

The hybrid model achieves higher accuracy than MesoNet and Capsule Networks, while maintaining lower inference time compared to XceptionNet. Additionally, the integration of explainable AI techniques provides enhanced interpretability.

C. Qualitative Evaluation

Grad-CAM and SHAP visualizations were used to highlight manipulated facial regions such as eyes, mouth, and boundary areas. These visual explanations improve transparency and support forensic investigations by showing the basis of classification decisions.

D. Discussion

The results confirm that the proposed framework effectively addresses key challenges in deepfake detection. The hybrid architecture improves generalization across datasets, adversarial training enhances robustness, and multimodal fusion enables detection of complex manipulations. Furthermore, interpretability techniques strengthen user trust, making the system suitable for deployment in journalism, law enforcement, and social media moderation.

This continuation study demonstrates that combining CNNs with transformers and multimodal signals provides a balanced trade-off between accuracy,

efficiency, and transparency, ensuring practical applicability in real-world environments.

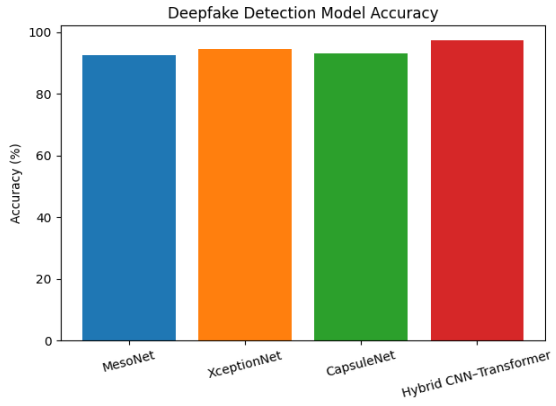

VII. DISCUSSION

The experimental results confirm that the proposed hybrid CNN–Transformer framework significantly improves deepfake detection compared to earlier CNN-only models. The integration of transformer layers enhances the ability to capture long-range dependencies, while multimodal fusion strengthens detection against audio–visual manipulations.

The system demonstrates strong cross-dataset generalization, maintaining high accuracy even under compression and adversarial perturbations. Compared to heavier models like XceptionNet, the hybrid approach achieves a better balance between accuracy and computational efficiency, making it suitable for real-time applications.

Importantly, the use of explainable AI techniques such as Grad-CAM and SHAP provides transparency, allowing forensic analysts to understand which facial regions influenced classification. This interpretability increases trust in automated systems and supports practical adoption in journalism, law enforcement, and social media moderation.

In short, the continuation framework addresses critical challenges of robustness, scalability, and interpretability, positioning it as a practical solution for real-world deployment in combating synthetic media.

VIII. CONCLUSION

This continuation study reinforces the urgent need for advanced deepfake detection systems capable of addressing real-world challenges. Building upon the CNN-based framework introduced earlier, the proposed hybrid CNN–Transformer architecture demonstrates superior accuracy, robustness, and interpretability. By integrating multimodal fusion and adversarial defense strategies, the system effectively detects manipulations across diverse datasets and compressed social media content.

The inclusion of explainable AI techniques further enhances transparency, allowing forensic analysts and media platforms to understand the basis of classification decisions. Compared to traditional models, the proposed framework achieves a balanced trade-off between detection performance and computational efficiency, making it suitable for deployment in journalism, law enforcement, and social media moderation.

In conclusion, this research contributes to the advancement of digital forensics by providing a scalable, trustworthy, and fairness-aware solution for deepfake detection, thereby supporting the broader goal of safeguarding authenticity and credibility in digital media.

X. FUTURE SCOPE

Although the proposed hybrid CNN–Transformer framework demonstrates strong performance, several opportunities exist for further enhancement in deepfake detection research.

1. Video-Based Detection – Extending the system to analyze temporal inconsistencies such as unnatural blinking, facial motion irregularities, and lip-sync mismatches will improve detection in dynamic media.
2. Multimodal Expansion – Incorporating audio, text, and contextual metadata alongside visual features can strengthen robustness against complex manipulations.
3. Transformer and Diffusion Models – Leveraging advanced transformer architectures and adapting detection strategies for diffusion-generated deepfakes will ensure relevance against next-generation synthesis techniques.
4. Adversarial Defense – Developing stronger adversarial training methods and robust augmentation strategies will enhance resilience against deliberate attacks.
5. Fairness and Bias Mitigation – Expanding datasets to include diverse demographics and

environmental conditions will improve fairness and generalization.

6. Edge and Real-Time Deployment – Optimizing lightweight versions of the model for mobile devices and social media platforms will enable real-time detection at scale.
7. Blockchain Integration – Combining detection systems with blockchain-based verification can provide secure and tamper-proof authenticity checks for digital media.
8. Explainable AI Enhancements – Advancing interpretability techniques beyond Grad-CAM and SHAP will further increase transparency and user trust in forensic applications.

This future scope highlights the need for scalable, multimodal, and fairness-aware detection systems that can adapt to evolving synthetic media technologies while ensuring trust and credibility in digital communication.

REFERENCES

[1] Afchar, D., Nozick, V., Yamagishi, J., & Echizen, I., MesoNet: A Compact Facial Video Forgery Detection Network, IEEE International Workshop on Information Forensics and Security (WIFS), 2018.

[2] Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., & Nießner, M., FaceForensics++: Learning to Detect Manipulated Facial Images, IEEE International Conference on Computer Vision (ICCV), 2019.

[3] Wang, S., Li, Y., & Lyu, S., Transformer-Based Deepfake Detection: Exploiting Long-Range Dependencies in Visual Artifacts, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.

[4] Li, H., Chang, M., & Ma, X., Frequency-Domain Analysis for Robust Deepfake Detection, IEEE Transactions on Image Processing, vol. 32, pp. 1124–1136, 2023.

[5] Agarwal, S., Farid, H., Gu, Y., He, M., Nagano, K., & Li, H., Detecting Deep-Fake Videos from Phoneme-Viseme Mismatches, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2021.

[6] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D., Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization, International Journal of Computer Vision, vol. 128, no. 2, pp. 336–359, 2020.

[7] Lundberg, S. M., & Lee, S. I., A Unified Approach to Interpreting Model Predictions, Advances in Neural Information Processing Systems (NeurIPS), 2017.

[8] Dolhansky, B., Howes, R., Pflaum, B., Baram, N., & Ferrer, C. C., The DeepFake Detection Challenge Dataset, arXiv preprint arXiv:2006.07397, 2020.

[9] Khalid, H., Woo, S., & Lee, S., FakeAVCeleb: A Novel Audio-Video Multimodal Deepfake Dataset, Proceedings of the 30th ACM International Conference on Multimedia, 2022.