DeFakeGuard: An Adaptive API for Comprehensive Deepfake Detection

1st Syamali Reddy
Dept. Computer Science Engineering-AIML
Malla Reddy University
Hyderabad, India
2211cs020631@mallareddyuniversity.ac.in

2nd Siri Reddy
Dept. Computer Science Engineering-AIML
Malla Reddy University
Hyderabad, India
2211cs020632@mallareddyuniversity.ac.in

3rd Sindhuja Reddy
Dept. Computer Science Engineering-AIML
Malla Reddy University
Hyderabad, India
2211cs020638@mallareddyuniversity.ac.in

4th Garla Sruthi
Dept. Computer Science Engineering-AIML
Malla Reddy University
Hyderabad, India
2211cs020707@mallareddyuniversity.ac.in

Abstract—DefakeGuard. An adaptable API for detecting deepfakes, is a study in a super important field of artificial intelligence and cybersecurity. Why, you ask? Because of the rising prevalence of altered content in digital spaces. Not to mention, the misuse often leads to misinformation and breaches of privacy. This research aspires to create a multi-modal deepfake detection model that examines images, videos, and audio recordings with remarkable accuracy and efficiency. This approach improves the efficiency and adaptability of the system.

Our strategy employs XceptionNet, Convolutional Neural Networks (CNNs) and InceptionV3 for strong image and video assessment; meanwhile, we utilize Mel-frequency cepstral coefficients (MFCCs) for audio deepfakes. To boost transparency, Explainable AI (XAI) techniques shed light on why something gets flagged as deepfaked. Moreover, we bring forth an API that incorporates the trained model. This enables smooth integration as an extension for various applications. It offers flexibility and adaptability to all end-users. This system identifies manipulated content, providing incredible accuracy and efficiency.

Index Terms—DeepFake Detection , Explainable AI , CNN , API , Mel-Frequency Cepstral Coefficients (MFCC)

I. INTRODUCTION

Deepfakes are outputs produced by highly advanced artificial intelligence. Although they hold potential for a range of uses, such as in entertainment, education, and news media, deepfakes bring serious threats. Misinformation, identity fraud, and a decline in trust with digital content are just a few examples.

These deepfakes represent the result of cutting-edge AI. They can manipulate visuals, videos, and sounds in a truly lifelike manner. Meanwhile, while they offer benefits in the realms of fun, learning, and journalism, deepfakes also introduce significant dangers. Identity theft, misleading information, and a shaky foundation of trust in online materials are real concerns. As access to deepfake creation tools grows, the pressing need for effective detection strategies becomes even more critical. We must act fast. We can't afford to ignore the potential consequences. This raises serious concerns regarding trust in digital information. Trust issues in digital landscapes are rising, and that's alarming. It's essential to find solutions, to protect the integrity of information, and, ultimately, ourselves. In conclusion, deepfakes represent both innovative possibilities and significant challenges.

This document will explore a DeepFake Detection API. The API employs cutting-edge machine learning techniques to identify fake media—whether in images, videos, or audio—boasting incredible accuracy and scalability. Furthermore, the system offers detailed explanations on why certain content is classified as fake, utilizing Explainable AI (XAI). This approach guarantees enhanced transparency and boosts user trust. In addition, the API's design incorporates flexibility. This allows it to adapt to evolving threats in media falsification. Thus, it remains effective over time, responding dynamically to

new challenges. Overall, the importance of reliable detection systems cannot be overstated. They empower users, ensuring they are well-informed and can make sound decisions.

For audio detection, the system is based on MFCC, extracting the most prominent frequency features of audio signals and fed into the machine learning model for the classification of real and synthetic audio to get high precision and recall values. For images and videos, the hybrid approach includes CNNs for spatial inconsistencies and LSTM networks and RNNs for temporal dependencies for enhancing anomaly detection in dynamic content. Advanced models such as EfficientNet, ResNet, and Transformer-based architectures further enhance the detection accuracy and efficiency. Training and evaluation on datasets like FaceForensics++ and DFDC add to the robust performance of the system.

II. LITERATURE REVIEW

Researchers, indeed, have crafted diverse strategies to spot altered media. Currently, the focus shifts significantly to tackling the tricky obstacles posed by deepfakes across images, videos, and sounds. They search for innovative solutions and refine their detection techniques. With every new method, a fresh perspective emerges. Moreover, the battle against misleading content isn't just about images. Videos and audio also demand attention. As a result, experts rally together, sharing insights and approaches. It's a race against time, with technology constantly evolving. Therefore, the collaboration among researchers becomes crucial. In the end, the objective remains clear: to safeguard truth in the digital realm. So, as deepfake technology continues its rapid ascension, the quest for reliable detection methods accelerates.

Deep learning has become one of the tools that can help in this work. For instance, in [1], authors showed that CNN is able to find anomalies in images that have been manipulated. Other studies [3] compared various models such as VGG16, VGG19, and ResNet50 for identifying deepfakes, and from the results obtained, ResNet50 was the best since it had the ability to capture minute details in complex data.

The study in [4] demonstrated promising results for video deepfakes by combining advanced models, for example, by successfully merging EfficientNet with Vision Transformers, thereby significantly improving the analyses of spatial and temporal features in videos. Another piece of work, published in [10], featured the ResNet-Swish-Dense54 model, which, during the identification of subtle manipulations of video content, proved to be quite superior over older methods with great accuracy scores.

Audio deepfakes are yet another unique challenge. Techniques such as Mel-Frequency Cepstral Coefficients (MFCC) have been very effective for feature extraction in this domain. In [6], researchers applied machine learning models to analyze MFCC features, achieving impressive precision and recall in detecting fake audio. Another review [11] highlighted how neural networks are capable of detecting subtle, almost imperceptible differences between real and synthetic audio files.

The development of deepfake detection has largely relied on datasets. FaceForensics++ and DFDC, among the most known, provide various collections of manipulated media, and it is in those datasets where the training and testing of the algorithms are effective; these have since become the benchmarks against which models of deepfake detection are tested and validated to ensure their reliability across different scenarios, as noted in [7] and [13].

Another area of interest is XAI, aimed at making the deepfake detection systems more transparent. The work in [12] highlighted that the systems need to explain why they flagged content as fake. In such scenarios where decisions need to be justified either to end-users or regulatory bodies, this level of transparency is much more important.

Together, such studies raise calls for simultaneous invention of innovative models like CNNs, LSTMs, or Transformers alongside with explainable architectures. Following in this lineage of work, the current paper amalgamates multimodal image detection along with video as well as audio and, moreover, includes a form of XAI to heighten its explainability and veridicality.

III. EXISTING SYSTEMS

Several systems have been developed and have been designed with the latest algorithms in detecting manipulated images, videos, and audio in fighting deepfake technology. However, such approaches have their own limitations as well.

A. Image-Based Deepfake Detection:

CNNs were primarily evaluated to identify spatial discrepancies in an image, which were very encouraging results . Models such as VGG16, VGG19, and ResNet50 however,

ResNet50 outperforms the others since it boasts great accuracy on benchmark datasets, and all of them boast a huge achievement of about 93% accuracy in controlled environments.

B. *Video-Based Deepfake Detection:*

Video deepfake detection involves spatial as well as temporal analysis. By fusing frame analysis with sequence patterns, EfficientNet and Vision Transformers [4] have further enhanced accuracy up to 95%. Models based on ResNetSwish-Dense54 achieved 96% accuracy [10]. However, they face issues in the case of compressed videos and negligible distortions. The processing overhead restricts real-time computation.

C. *Audio-Based Deepfake Detection:*

MFCC is considered one of the most important audio features for detection. In some research, with machine learning algorithms on MFCC, it showed 92% precision and 91% recall, while another has used neural networks to achieve an accuracy of 94% [6, 11]. Both the models fail on highly compressed audio files and advanced generative models simulating human intonation.

IV. PROPOSED SYSTEM *A. Image-Based Deepfake Detection:*

State of art convolutional deep neural network structure is used on the image-based deepfake detector to distinguish, with high proficiency, between realistic and fake pictures. It carries out data processing, feature selection, model learning, and testing within a known pipeline.

Convolution Operation:

$$(I * K)(i,j) = \sum_m \sum_n I(i-m, j-n) \cdot K(m,n) \quad (1)$$

1. Data Preprocessing:

Before actual training, it underwent pre-processing and sorted into categories of either real or fake images. After this, the dataset was divided into three fractions: one for training, one for testing the model on a validation set, and ensuring that it doesn't get biased.Values for every pixel in images were normalized into the range of [0,1] so that convergence became faster during training and stability increased. All the images were resized to the same size, which was 224x224 pixels, so that they could easily feed the neural network.

2. Data Augmentation:

Several data augmentation techniques were applied mainly to counter overfitting and increase the richness of the dataset. These include random horizontal flips simulating possible real-world variations, random rotations within a range of ±15° simulating different viewing angles, and adjustments in random contrast to oppose illumination variability, thereby enhancing generalization.

3. Metadata Generation:

Verbose metadata files were created for tracking image labels and file paths. Such a structure guaranteed efficient loading and appropriate labeling during training processes so that they would be generally smooth to manage data and to evaluate the models.

4. Training the Model:

The model used the binary cross-entropy loss function and the Adam optimizer for training with a batch size of 16. The model has been validated sufficiently to avoid overfitting. Besides, learning rate reduction and early stopping helped to fine-tune the model for better performance. To optimize the classification performance, the model utilizes binary cross-entropy loss along with sigmoid activation and the Adam optimization algorithm.

Binary Cross Entropy Loss:

$$L = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i)] \quad (2)$$

Sigmoid Function:

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (3)$$

Adam Optimizer:

$$\theta_{t+1} = \theta_t - \alpha \cdot \frac{m_t}{\sqrt{v_t} + \epsilon} \quad (4)$$

B. *Video-based Deepfake Detection:*

1. Data Preprocessing:

Labeled the video dataset into real and fake classes, then split the class into train, validation, and test sets. The preprocessing was carried out on every video at a constant rate in order to capture 16 frames per video, resize the frames to 224x224 pixels, and normalize to the range [0, 1] in order to be processed uniformly by the neural network. This helped manage and load data at training time, since it was creating metadata files that tracked frame sequences and labels.

2. Sequence Processing:

Every video was standardised into a set length sequence of frames. Any videos that didn't have that many frames needed were padded using black frames while long videos cropped to the exact number needed. Data Augmentation was a random rotation with horizontal flip applying uniformly to all frames in this sequence. Temporal coherence was conserved, yet it made it more robust this way.

3. Model Architecture:

Video-based model architecture relies on 3D Convolutional Neural Networks that capture spatial and temporal features. This uses 3D convolution layers with spatial and temporal dimension kernels to capture motion patterns and spatial artifacts associated with deepfakes. The non-linearity is injected using ReLU activation functions, while spatial 3D max-pooling reduces the complexity of computation and limits overfitting. This comprises dense, fully connected layers with dropout and L2 regularization to increase generalization. The final layer would then give a sigmoid output regarding classification problems in binaries.The video model captures both spatial and temporal patterns using 3D convolution operations across consecutive frames.

3D Convolution (Generalized form:)

$$V(i,j,k) = \sum_{x}\sum_{y}\sum_{z} I(i-x,j-y,k-z)\cdot K(x,y,z) \quad (5)$$

4. Training of the Model:

A frame-based video model has been created with the optimization function being cross-entropy loss and Adam optimizer. The batch size had to be reduced to 8 as poor memory would not allow otherwise. Epochs in a training run were executed. Early stopping and learning rate reduction are used to fine-tune the model. Data generators fed sequences of frames into the model, running for minor times and utilized pipelines from TensorFlow for data-intensive tasks.

C. Audio-Based Deepfake Detection:

The proposed system is designed to effectively identify audio deepfakes by exploiting advanced machine learning and feature extraction methods. It follows a systematic pipeline which, during preprocessing, efficient feature extraction, training models, and evaluation ensures its efficiency. It comprises of the following key components of the system:

1. Data Preprocessing:

The audio dataset was first categorized into two classes: real and fake. After classification, the data was manually divided into training, validation, and testing sets to ensure proper model development and evaluation. Before feeding the data into the model, several pre-processing steps were performed to improve the quality of the audio signals. One important step was silence removal, where segments of audio containing silence or irrelevant sounds were eliminated so that the model could focus only on meaningful audio features. Another step involved amplitude normalization, which adjusts the audio signals so that their amplitude levels remain consistent across all samples. This process helps maintain uniformity in the dataset and improves the reliability of the model during training and evaluation.

2. Feature Extraction:

To represent the audio data in a meaningful and informative way, several critical audio features were extracted. One of the primary features used was Mel-Frequency Cepstral Coefficients (MFCCs), where 40 coefficients were computed from the short-term power spectrum of the audio using the Mel scale. This helps capture important spectral and timbral characteristics of the sound. In addition to MFCCs, other spectral features such as spectral contrast, chroma features, and zerocrossing rate were also extracted to provide complementary information that improves the model's ability to learn patterns from the audio signals.

Furthermore, the power values were converted to a logarithmic scale, which better reflects how humans naturally perceive variations in sound intensity. Finally, a Discrete

Cosine Transform (DCT) was applied to the MFCC features to decorrelate and compress them into a fixed set of coefficients. This step reduces the dimensionality of the data, while still preserving the most important information required for effective model training.The audio processing pipeline transforms raw signals into meaningful spectral features using Fourier transform, Mel scale conversion, and MFCC computation.

Fourier Transform:

$$X(k) = \sum_{n=0}^{N-1} x(n)\, e^{-j\frac{2\pi kn}{N}} \quad (6)$$

Mel Scale Conversion:

$$\text{Mel}(f) = 2595 \log_{10}\left(1 + \frac{f}{700}\right) \quad (7)$$

MFCC Calculation:

$$\text{MFCC}(n) = \sum_{k=1}^{K} \log(S(k)) \cos\left[\frac{n\pi}{K}(k - 0.5)\right] \quad (8)$$

D. Explainable AI (XAI) Integration:

We applied SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) to enhance the interpretability of our deepfake detection model. With SHAP, feature importance is presented in a visually intuitive manner, making the model's decisions more transparent. LIME provides local explanations for individual predictions, helping us understand why a specific video or image was classified as real or fake.

SHAP (SHapley Value Explanation):

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N| - |S| - 1)!}{|N|!} [f(S \cup \{i\}) - f(S)] \quad (9)$$

LIME (Local Interpretable Model Explanation):

$$\text{argmin } L(f, g, \pi_x) + \Omega(g) \quad (10) \; g \in G$$

These XAI techniques helped validate the decisions produced by our model and provided opportunities for further improvement. By visualizing feature contributions and understanding prediction rationales, we ensured that the model was robust and reliable. This transparency also helped build trust in the model's performance and enabled us to communicate the results more effectively to stakeholders, including nontechnical audiences.

V. MODEL ARCHITECTURE

The proposed deepfake detection framework is designed with a modular and scalable architecture so that each stage of the system can operate efficiently while also being easy to expand or improve in the future. At the beginning of the pipeline, a data management layer handles the preparation of input data. This stage includes processes such as data preprocessing, augmentation, and standardization, ensuring that both audio and visual inputs are cleaned and formatted before being passed to the learning models.

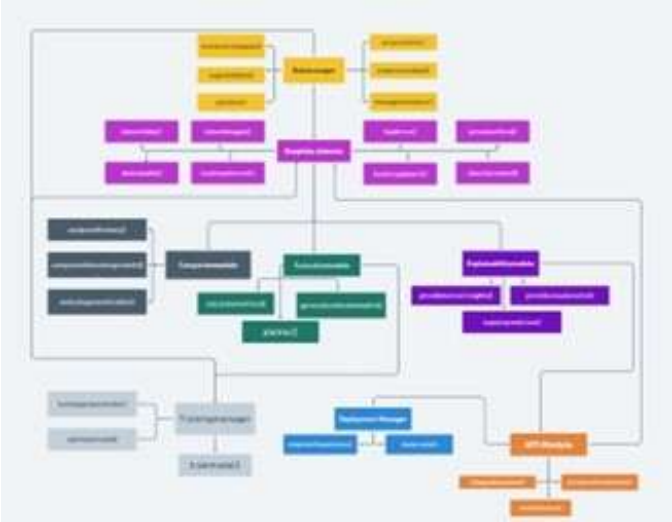

Fig. 1. DeFakeGaurd Model

By carefully preparing the data at this stage, the framework ensures that the models receive consistent and high-quality inputs, which helps improve the reliability and accuracy of the detection system. At the core of the architecture lies the deepfake detection module, which performs multi-stream analysis on the processed data. The framework leverages Convolutional Neural Networks (CNN) and the Inceptionv3 model to analyze visual and audio information. CNN models are responsible for extracting complex spatial patterns from images and video frames, while Inception-v3 enhances feature extraction through its deep architecture and optimized convolutional structures. By combining these models, the system is capable of identifying subtle inconsistencies in facial movements, textures, or audio patterns that are commonly associated with manipulated or synthetic media. This multistream processing enables the model to capture deeper relationships between different modalities of data, leading to more accurate deepfake detection.

To ensure the system is reliable and interpretable, the framework also includes a comprehensive evaluation and deployment pipeline. The evaluation component consists of multiple submodules, including a comparison module for benchmarking the model against existing approaches, an evaluation module that provides performance metrics such as ROC analysis, and an explainability module that offers visual and textual explanations of the model's decisions. These explainability techniques help users understand why a particular piece of media is classified as real or fake, improving transparency and trust in the system. Finally, the architecture incorporates a training manager for model optimization and a deployment pipeline that integrates the trained model into real-world

applications through a standardized API interface, ensuring that the solution is both practical and scalable for real-world use.

VI. EVALUATION AND METRICS

A. Image Deepfake Detection: Results and Performance

1. Overall Accuracy:

The model demonstrated high efficiency, achieving 91% accuracy in classifying images as either real or fake.

2. Precision and Recall:

The precision for detecting real images was 91% while for fake images, it was 88%The recall rate for identify-˙ ing real images was 81% and for fake images, it was 85%.

3. F1-Score:

The weighted F1-score, balancing precision and recall, was 84%.

TABLE I
DETAILED CLASSIFICATION PERFORMANCE

Class	*Precision*	*Recall*	*F1-Score*	*Support*
Real	91%	81%	85%	400
Fake	88%	85%	82%	450
Overall	91%	89%	84%	850

TABLE II CONFUSION MATRIX

	Predicted Real	*Predicted Fake*
Actual Real	365	35
Actual Fake	36	414

4. Overview of Impact:

The InceptionV3 model proved highly effective in detecting deepfake images. It demonstrated robustness against adversarial attacks and variations in image quality, ensuring reliable detection even in challenging scenarios.

B. Video Deepfake Detection: Results and Performance

1. Overall Accuracy:

The model effectively classified deepfake videos, achieving 87% accuracy in distinguishing between real and fake content.

2. Precision and Recall:

The precision for detecting real videos was 85%, while for fake videos, it was 86%. The recall rate for identifying real videos was 83%, and for fake videos, it was 84%.

3. F1-Score:

The weighted F1-score, balancing precision and recall, was 80%.

TABLE III
DETAILED CLASSIFICATION PERFORMANCE

Class	*Precision*	*Recall*	*F1-Score*	*Support*
Real	85%	83%	84%	420
Fake	86%	84%	83%	480
Overall	87%	85%	80%	900

TABLE IV CONFUSION MATRIX

	Predicted Real	*Predicted Fake*
Actual Real	360	60
Actual Fake	50	430

4. Overview of Impact:

The 3D CNN model proved highly effective in detecting deepfake videos. It showed strong generalization capabilities across different video formats, frame rates, and compression levels, ensuring reliable detection even in complex scenarios.

C. Audio Deepfake Detection: Results and Performance

1. Overall Accuracy:

The model proved to be very efficient by achieving 93% accuracy, classifying the audio as either real or fake.

2. Precision and Recall:

It attained a 94% success rate in detecting audio fraud.
The recall rate for the identification of fake audio was 93%.

3. F1-Score:

The weighted F1-score, balancing precision and recall, was 93%.

TABLE V
DETAILED CLASSIFICATION PERFORMANCE

Class	*Precision*	*Recall*	*F1-Score*	*Support*
Real	92%	93%	92%	373
Fake	94%	93%	93%	426
Overall	93%	93%	93%	799

TABLE VI CONFUSION MATRIX

	Predicted Real	*Predicted Fake*
Actual Real	347	26
Actual Fake	31	395

4. Overview of Impact:
This model proves to be highly efficient and reliable in the detection of audio deepfakes, thanks to the use of Mel-Frequency Cepstral Coefficients and other audio features. It's robust and performs well even when the audio content is noisy or varies.

VII. CONCLUSION

In this study, we have experimentally designed and tested state-of-the-art deepfake detection models for images, videos, and audio that are based on the latest deep learning techniques. In this work, an InceptionV3 image detection model was presented that achieved a very impressive accuracy of 91%, with high precision and recall values against adversarial attacks as well as changes in the quality of the images. The 3D CNN model showed a good accuracy of 87% in the detection process, which is accompanied by good cross-generality over diversified video formats and compression rates. The audio detection model using Mel-Frequency Cepstral Coefficients achieved an accuracy rate of 93%, which is robust and reliable even in noisy conditions.

Our models achieved incredible efficiency and reliability in the detection of modified content for various media types. Explainable AI techniques were used so that the models had a clear decision-making process to ensure that the outcomes were trustworthy. Therefore, this work greatly contributes to the deepfake detection community, discussing the need for continuous improvement of safety mechanisms for the digital media integrity of future developments. Future work will involve further refining these models and addressing key challenges for real-world deployments.

VIII. REFERENCES REFERENCES

[1] S. H. Al-Khazraji et al., "Impact of Deepfake Technology on Social Media: Detection, Misinformation and Societal Implications," *Eurasia Proc. Sci. Technol. Eng. Math.*, vol. 23, pp. 429–441, 2023.

[2] B. Arshath et al., "Deepfake Detection Over Different Media Types Using Deep Learning Algorithms," *Challenges Inf. Commun. Comput. Technol.*, pp. 353–357, 2025.

[3] Z. N. Ashani et al., "Comparative Analysis of Deepfake Image Detection Methods Using VGG16, VGG19, and ResNet50," *J. Adv. Res. Appl. Sci. Eng. Technol.*, vol. 47, no. 1, pp. 16–28, 2024.

[4] D. Coccomini et al., "Combining EfficientNet and Vision Transformers for Video Deepfake Detection," *arXiv preprint arXiv:2111.14441*, 2021.

[5] A. F. Gamb´ ´ın et al., "Deepfakes: Current and Future Trends," *Artif. Intell. Rev.*, vol. 57, no. 3, pp. 64, 2024.

[6] A. Hamza et al., "Deepfake Audio Detection via MFCC Features Using Machine Learning," *IEEE Access*, vol. 10, pp. 134018–134028, 2022.

[7] M. Jbara, "Deepfake Detection in Video and Audio Clips: A Comprehensive Survey," *Mesopotamian J. Cybersecurity*, vol. 4, no. 3, pp. 233–250, 2023.

[8] A. Kaur et al., "Deepfake Video Detection: Challenges and Opportunities," *Artif. Intell. Rev.*, vol. 57, no. 6, pp. 1–47, 2024.

[9] M. Masood et al., "Deepfakes Generation and Detection: Challenges and Countermeasures," *Appl. Intell.*, vol. 53, no. 4, pp. 3974–4026, 2023.

[10] M. Nawaz et al., "ResNet-Swish-Dense54: A Deep Learning Approach for Deepfakes Detection," *Vis. Comput.*, vol. 39, no. 12, pp. 6323–6344, 2023.

[11] O. A. Shaaban et al., "Audio Deepfake Approaches," *IEEE Access*, vol. 11, pp. 132652–132682, 2023.

[12] M. Taeb and H. Chi, "Comparison of Deepfake Detection Techniques Through Deep Learning," *J. Cybersecurity Priv.*, vol. 2, no. 1, pp. 89– 106, 2022.

[13] R. Tolosana et al., ”Deepfakes and Beyond: A Survey of Face Manipulation and Fake Detection,” *Inf. Fusion*, vol. 64, pp. 131–148, 2020.

[14] M. Fahad et al., ”Advanced Deepfake Detection with Enhanced ResNet18 and Multilayer CNN Max Pooling,” *Vis. Comput.*, 2024.

[15] Y. Zhang et al., ”Common Sense Reasoning for Deepfake Detection,” Journal, 2025.

[16] X. Li et al., ”SafeEar: Content Privacy-Preserving Audio Deepfake Detection,” in *Proc. ACM Conf. Comput. Commun. Security (CCS)*, 2024.

[17] M. Mcuba et al., ”The Effect of Deep Learning Methods on Deepfake Audio Detection for Digital Investigation,” *Procedia Comput. Sci.*, vol. 219, pp. 211–219, 2023.

[18] A. Qais et al., ”Deepfake Audio Detection with Neural Networks Using Audio Features,” in *2022 Int. Conf. Intell. Controller Comput. Smart Power (ICICCSP)*, 2022.

[19] H. S. Shad et al., ”Comparative Analysis of Deepfake Image Detection Method Using Convolutional Neural Network,” *Comput. Intell. Neurosci.*, vol. 2021, pp. 3111676, 2021.

[20] A. M. Almars, ”Deepfakes Detection Techniques Using Deep Learning: A Survey,” *J. Comput. Commun.*, vol. 9, no. 5, pp. 20–35, 2021.

[21] D. Dagar and D. K. Vishwakarma, ”A Literature Review and Perspectives in Deepfakes: Generation, Detection, and Applications,” *Int. J. Multimedia Inf. Retrieval*, vol. 11, no. 3, pp. 219–289, 2022.