



DEEPFAKE DETECTION USING CONVOLUTIONAL NEURAL NETWORKS: A SECURE MEDIA AUTHENTICATION SYSTEM

A Deep Learning Approach for Image and Video Forgery Detection

¹Chinna Siva Krishna Thota, ²Sarika Dasari.

¹M.Tech, ²Assistant Professor

¹Department of Computer Science & Engineering, ²Department of Computer Science & Engineering

¹Bhimavaram Institute of Engineering and Technology, Bhimavaram, India, ²Bhimavaram Institute of Engineering and Technology, Bhimavaram, India

Abstract : Deep fake technology has developed significantly, raising important issues with the authenticity of digital media, privacy, and security. In this study, a system for detecting manipulated images and videos is proposed based on the Convolutional Neural Network (CNN). The CNN is trained with a dataset that includes both original and manipulated images and videos. The performance of the proposed system is evaluated based on standard performance measures. The results of the experiment show that the proposed system can achieve 98% accuracy in image classification and can analyze the frames of the video. The proposed system is developed based on the Django web technology to provide an easy way to authenticate the digital media. The proposed system is compared with other existing systems to show the efficiency of the proposed approach.

IndexTerms: Deepfake Detection, Convolutional Neural Networks, Media Authentication, Video Forensics, Computer Vision.

I. INTRODUCTION

Deepfake technology uses intelligence, especially something called Generative Adversarial Networks to make fake media that looks really real. This kind of content can be used to do bad things like spread false information steal peoples identities and cheat others. Deepfake technology can be used for these activities, which is why we need good systems to detect it. The old ways of detecting content often have trouble finding fake things made by artificial intelligence so we need new systems that can learn and get better. Deepfake technology is a problem because it can make realistic fake media. The deepfake methods are getting better and better. This is a problem for people who try to figure out if a video or picture is real or not. It affects a lot of areas like the news, police work and keeping peoples information safe. The deepfake methods are a concern for digital forensics and media verification and they have a big impact on things, like journalism, law enforcement, cybersecurity and individual privacy.

Over the few years the way we make deepfakes has gotten a lot better. This is because of technology like StyleGAN, Progressive GAN and Variational Autoencoders. Now deepfakes can copy the way someone looks, sounds and moves. It is really hard to tell what is real and what is not. This is a problem because people are worried that deepfakes will be used to trick people like spreading false information about politicians cheating people out of money hurting someones reputation and making people not trust what they see on the internet on things like Facebook and YouTube and other places where people share videos and pictures. Deepfakes are getting better and better. That is why people are so worried, about them.

This study is about a system that can detect deepfakes. The system uses computer technology to find out if a picture or video is real or not. It is built into a website that people can use easily. Anyone can upload a picture or video to the website to see if it is real.

The people in charge of the website can look at a lot of data change how the system works and make it better over time. They can also see how well the system is working and look at the results, in a way. The goal of this system is to help people figure out if a picture or video has been manipulated. Deepfake detection is very important. This system is trying to meet that need. It wants to be a tool that people can use to detect manipulated media like deepfakes.

The main contributions of this work include a highly precise CNN-based model that achieves 98.47% accuracy in distinguishing real images from fake ones; a frame-level video analysis mechanism that checks temporal consistency and motion patterns to detect manipulated footage; a secure and scalable web system with role-based access for users and administrators; thorough performance evaluation using metrics such as accuracy, precision, recall, F1-score, and AUC-ROC across varied testing conditions; comparative testing that demonstrates improved results over several established detection models; and detailed ablation experiments that analyze how different architectural elements and preprocessing steps affect overall performance.

II. RELATED WORK

In recent years, deepfake detection has become a highly active research area due to the rapid advancement of media generation techniques. Researchers have investigated multiple strategies for identifying manipulated content using machine learning, deep learning, and computer vision methods. For example, Rossler and colleagues introduced FaceForensics++, a large-scale benchmark dataset with more than 1.8 million altered video frames designed for facial forgery detection. This dataset includes several manipulation approaches such as Face2Face, FaceSwap, DeepFakes, and NeuralTextures, and is widely used as a standard benchmark for evaluating detection models in academic studies.

Afchar and collaborators developed MesoNet, a lightweight convolutional neural network tailored for detecting forged videos. Their method focuses on mesoscopic image characteristics—features that lie between fine details and broader structures—which are often distorted during deepfake creation. Despite its compact design with only four convolutional layers, MesoNet achieved an accuracy of 95.3% on the FaceForensics dataset, demonstrating that efficient architectures can still produce strong detection results.

Transfer learning has also played an important role in this field. Li and co-authors showed that analyzing facial warping artifacts can serve as an effective detection signal, reaching 97.3% accuracy on FaceForensics++. Their approach identifies irregularities in facial landmark alignment caused by affine transformations commonly applied in face-swapping techniques. Similarly, Nguyen and team proposed a capsule network-based model that examines spatial relationships among facial features and remains resilient to common image modifications such as compression, blurring, and noise.

Another key factor in identifying manipulated videos is temporal consistency. Güera and Delp applied Long Short-Term Memory (LSTM) networks to study temporal patterns in video sequences, significantly outperforming methods that analyze frames individually. Their model detects inconsistencies across frames in expressions and head movements that often appear in synthetic videos. Building on this concept, Sabir and colleagues introduced recurrent convolutional approaches to capture subtle irregularities in facial muscle motion and blinking behavior.

Attention mechanisms have further enhanced detection performance. Zhao and associates designed multi-attention networks that concentrate on critical facial regions such as the eyes, nose, mouth, and edges of the face. These modules learn which areas provide the most useful information for identifying manipulation. In a related effort, Kumar and collaborators integrated both spatial and channel attention, enabling models to focus simultaneously on meaningful locations and relevant feature maps.

Although these developments have greatly improved detection accuracy, many current solutions still encounter obstacles when applied in real-world environments. Common issues include limited accessibility for non-technical users, absence of real-time processing capabilities, susceptibility to adversarial manipulation, difficulty generalizing to unseen deepfake techniques, high computational demands that restrict deployment on low-resource devices, and a narrow focus on facial deepfakes rather than other manipulated media types such as synthetic speech or full-body alterations.

To address these limitations, this work introduces a hybrid detection framework that merges CNN-based image analysis with advanced frame-level video evaluation and deploys it within a scalable, production-ready web platform. The system closes the gap between research and practical use by delivering a complete solution featuring an intuitive interface, detailed performance monitoring tools, expandable architecture for future upgrades, strong security measures such as authentication and input validation, and thorough documentation for both administrators and developers.

III. SYSTEM ARCHITECTURE

The proposed system architecture is designed according to modern software engineering practices, using a modular structure that promotes scalability, maintainability, flexibility, and reliability. It follows a multi-layered architecture consisting of presentation, application, and data layers, each responsible for specific functions. This structured approach allows individual components to be developed, tested, deployed, and scaled independently, supporting agile workflows and continuous integration. Such modularity also ensures that updates or improvements can be made to one part of the system without disrupting the entire platform.

The solution operates on a client-server model in which a responsive frontend interface, built with HTML5, CSS3, and JavaScript, communicates with a Django-based backend through clearly defined RESTful APIs. The backend manages essential operations such as user authentication, file uploads, preprocessing tasks, model inference requests, result formatting, and database storage. The detection engine itself functions as a microservice that can be scaled across multiple GPU-enabled servers depending on system demand and workload, ensuring consistent performance even under heavy usage.

A. User Module

The user module is designed to provide a smooth and intuitive experience for end-users interacting with the system. It incorporates secure authentication mechanisms and a wide range of features. Users can register accounts through email verification with token-based confirmation and strong password requirements. The login system includes session management, protection against cross-site request forgery, automatic session timeout after inactivity, and safeguards that temporarily lock accounts after repeated failed login attempts. The platform supports uploading multiple media formats, including common image and video types, with strict validation checks for file type, size limits, and security scanning. Detection results are displayed instantly with confidence scores and color-coded indicators, along with probability details and explanations when available. Users can also access a history

of previous analyses, filter or sort records, compare results, and download past reports. In addition, the system can automatically generate detailed PDF reports containing analysis outcomes, visual summaries, timestamps, metadata, and recommendations when needed.

B. Admin Module

The administrative module equips system managers with comprehensive tools to oversee, configure, and optimize the platform. An interactive dashboard provides real-time statistics such as total detections, active users, resource utilization, performance trends, and alerts for unusual activity. Administrators can manage accounts, assign roles, monitor user actions, suspend accounts if necessary, and perform bulk operations efficiently. The module also includes dataset management capabilities for uploading training data, organizing files, labeling content, configuring augmentation settings, and maintaining dataset versions for reproducibility. Training tools allow administrators to adjust hyperparameters, select optimization methods, distribute training across GPUs, monitor progress with visual graphs, save checkpoints automatically, and stop training early if validation results plateau. Evaluation utilities generate confusion matrices, ROC curves, precision–recall charts, per-class metrics, and cross-validation reports. Advanced visualization features present performance analytics through interactive charts and statistical summaries, while maintenance tools support logging, error monitoring, backups, model version control, and automated health diagnostics.

C. Detection Module

The detection module serves as the core intelligence of the system, implementing advanced deep learning techniques to identify manipulated media accurately. It includes a preprocessing pipeline that standardizes inputs by resizing images, normalizing pixel values, detecting faces, and extracting relevant regions. Additional augmentation methods such as rotation, flipping, brightness and contrast adjustments, noise addition, and color variation help improve the model's ability to generalize across diverse inputs. The classification architecture uses a MobileNetV2 backbone pre-trained on a large dataset for feature extraction, followed by custom layers for classification. These layers include global average pooling, normalization components, dense layers for hierarchical feature learning, dropout regularization to reduce overfitting, and a final softmax layer that produces probabilities for real or fake categories.

The system's classification engine provides predictions along with confidence scores and adjustable thresholds that can be tuned depending on whether precision or recall is prioritized. It can also combine outputs from multiple model versions to increase reliability and includes uncertainty estimation to assist decision-making. Post-processing features enhance interpretability and robustness, such as heatmaps that highlight influential regions in an image, temporal consistency analysis for videos, anomaly detection for unfamiliar manipulation patterns, and explainable outputs that clarify the reasoning behind each prediction.

IV. METHODOLOGY

A. Dataset Preparation

The dataset for this study was assembled through a careful selection process that drew from multiple trusted sources to ensure diversity, balance, and realistic representation. The primary data came from the widely recognized FaceForensics++ benchmark, which includes both original videos and manipulated versions generated using advanced techniques such as facial reenactment, face swapping, GAN-based synthesis, and texture transfer. To strengthen the dataset further, additional authentic images were incorporated from Labeled Faces in the Wild, which contains natural facial variations from over 13,000 individuals, as well as the CelebA dataset, known for its wide range of celebrity faces and attributes. Custom-collected images representing different demographic backgrounds were also added to support fairness and reduce bias in model predictions.

The finalized dataset consisted of 2,041 training images and 2,041 validation images, evenly distributed between real and fake categories to avoid class imbalance. Each image was organized according to several criteria, including the manipulation method used, compression level, image quality characteristics such as resolution or noise, and demographic attributes like age, gender, and ethnicity. In addition, an independent test set of 814 images was created from entirely separate sources not involved in training or validation, enabling a more reliable assessment of how well the model performs on unfamiliar data.

To improve the model's ability to handle real-world variations and reduce overfitting, a wide range of augmentation techniques were applied. Geometric transformations included random rotations to simulate camera angle changes, horizontal flips to increase sample diversity, random cropping followed by resizing to encourage spatial robustness, and mild perspective adjustments. Photometric enhancements introduced changes in brightness, contrast, and color properties, along with simulated sensor noise and compression effects, helping the system adapt to different lighting conditions, image qualities, and device outputs.

Strict quality assurance procedures were implemented to maintain dataset reliability. Random samples from each category were manually reviewed by multiple annotators to confirm label accuracy, detect possible errors, evaluate visual clarity, and eliminate ambiguous or damaged images. Statistical checks were also conducted to verify balanced class distribution, ensure fair representation of manipulation types, confirm demographic diversity, and validate that the dataset included examples ranging from easily detectable alterations to more subtle and challenging manipulations.




Figure 1: Dataset Distribution Across Training and Validation Sets

B. Model Development

The CNN architecture used in this system is built around MobileNetV2, an advanced neural network known for combining strong performance with computational efficiency. This choice was guided by several advantages, including its high accuracy relative to parameter count, suitability for devices with limited resources, and proven success on large-scale image classification tasks. Its design features, such as inverted residual blocks, linear bottlenecks, and depthwise separable convolutions, help preserve important information while significantly reducing processing cost. Because of these strengths, MobileNetV2 has been widely adopted in real-world computer vision applications.

To make the most of this architecture, transfer learning was applied by initializing the base network with weights pre-trained on a large image dataset containing millions of labeled samples. This starting point provided the model with strong foundational feature recognition abilities, such as detecting edges, textures, and shapes, which are highly relevant for identifying facial manipulations. Training followed a gradual fine-tuning strategy: initially, only the custom classification layers were trained while the base network remained frozen; later, the top layers were unfrozen for refinement; and finally, nearly all layers were trained together for full optimization while keeping normalization layers fixed to maintain stability.

The final model structure begins with an input layer that accepts normalized 224×224 RGB images, followed by the MobileNetV2 backbone containing dozens of convolutional layers organized into bottleneck blocks with residual connections. A global average pooling layer converts spatial feature maps into a compact feature vector, which is then stabilized through batch normalization. The classification head consists of dense layers that learn higher-level patterns, with dropout layers included to reduce overfitting. The output layer uses a softmax function to produce probability scores indicating whether an image is real or fake.

The training procedure was carefully planned to achieve reliable convergence without overfitting. The Adam optimizer was chosen because of its adaptive learning rate and stability, and a dynamic learning rate scheduler automatically reduced the learning rate when validation performance plateaued. An early stopping mechanism monitored validation accuracy and halted training once improvement stopped, restoring the best-performing model weights. Training was conducted for up to 50 epochs using mini-batches of 32 samples, balancing computational efficiency with stable gradient estimation. Mixed-precision training was used on compatible GPUs to speed up calculations while maintaining numerical stability, and checkpoints were saved whenever validation performance improved. In practice, training required several hours on a high-performance GPU, processing dozens of samples per second.

C. Performance Metrics

Model evaluation relied on multiple metrics to provide a well-rounded understanding of performance. Accuracy measures overall correctness by comparing all correct predictions to total predictions, but it can be misleading when class distributions are uneven. Precision evaluates how reliable positive predictions are, indicating how often items labeled as fake are actually fake, which is especially important in contexts where false accusations could have serious consequences. Recall measures how effectively the system identifies all manipulated media, making it critical for applications where missing a deepfake could lead to significant harm.

The F1-score combines precision and recall into a single value using their harmonic mean, ensuring that the model performs consistently across both metrics rather than excelling in one while performing poorly in the other. Additional indicators further describe model behavior. Specificity shows how well authentic content is correctly recognized, reducing false alarms. The area under the ROC curve assesses how effectively the model distinguishes between classes across different decision thresholds. Another metric, the Matthews Correlation Coefficient, evaluates agreement between predictions and actual labels and is particularly useful when dealing with imbalanced datasets.

D. Video Detection Strategy

Detecting deepfakes in video is more complex than analyzing individual images because it requires consideration of temporal patterns, processing efficiency, and intelligent frame selection. To address these challenges, the system uses a multi-stage pipeline designed to balance accuracy with practical performance. The process begins by extracting frames at a standard rate to capture motion information while avoiding unnecessary duplication. For longer videos, an adaptive sampling method adjusts frame selection according to motion intensity, capturing more frames during rapid movement and fewer during static segments to conserve resources.

Each frame is then processed using multiple face detection techniques working together for reliability. These methods complement one another, ensuring accurate detection even under challenging conditions such as unusual angles, partial occlusion, or varying lighting. When multiple faces appear, the system selects the most confident detection result. Extracted faces are preprocessed using the same steps applied during training, including resizing, normalization, and optional enhancements to improve clarity in low-quality frames.

Every processed face is analyzed by the trained CNN model to generate predictions. These frame-level results are combined using a weighted voting approach that accounts for confidence levels, consistency across neighboring frames, and detection stability. Outlier predictions that appear unreliable are reduced in influence, and smoothing techniques help minimize noise in prediction sequences.

The final decision for each video is determined through a layered evaluation process that considers the majority of predictions, the stability of results over time, and the strength of evidence in individual frames. The system also highlights frames with the strongest indicators of manipulation for human review. Based on this analysis, videos are categorized as real, fake, or uncertain, with the last category indicating cases where automated analysis is inconclusive and further inspection is recommended.




Figure 2: Video Detection Accuracy Across Different Video Types

A range of performance optimization methods were implemented to support near real-time analysis. Frame batches are processed in parallel across multiple CPU cores to increase throughput, while GPU acceleration allows the CNN model to evaluate groups of 16–32 frames at once, significantly reducing inference time. The system also uses smart caching to store intermediate results so that similar frames do not need to be processed repeatedly, saving computational resources. In addition, a progressive streaming approach generates early predictions while the remaining video is still being analyzed, which is especially useful for longer clips. Together, these strategies enable processing speeds close to 1.3 times the video's length for standard-quality footage, making the system suitable for applications where timely results are essential.

V. EXPERIMENTAL RESULTS

Extensive experiments were carried out to thoroughly assess the effectiveness of the proposed deepfake detection system from several perspectives, including accuracy, processing speed, robustness to distortions, and comparison with well-known baseline models. All tests were conducted on a high-performance workstation equipped with an NVIDIA RTX 3080 GPU, 32GB of RAM, an Intel Core i7 processor, and fast NVMe storage to ensure smooth computation and rapid data handling. The evaluation process followed strict scientific standards so that results would be reliable, reproducible, and fairly comparable with existing approaches.

A. Image Classification Performance

The model demonstrated outstanding performance, achieving an accuracy of 98.47% on a separate test dataset of 2,414 images that were not used during training or validation. This result indicates strong generalization ability when handling unfamiliar data. The system outperformed several established architectures, including a baseline CNN model, MesoNet, ResNet-50, and XceptionNet, highlighting its effectiveness on this benchmark. A confusion matrix was also generated to provide a detailed breakdown of prediction outcomes for both real and fake categories, offering deeper insight into the model's classification behavior.

Actual \ Predicted	Real	Fake
Real	1578	22
Fake	15	799

Table 1: Confusion Matrix for Image Classification on Test Set

An examination of the confusion matrix shows that the model performs very well in identifying both real and manipulated content. Out of 1,600 genuine images, it correctly recognized 1,578, resulting in a specificity of 98.6%, which highlights its strong capability to detect authentic media accurately. In the case of altered content, the system successfully identified 799 of 814 fake images, achieving a recall rate of 98.2% and demonstrating high sensitivity to manipulated material.

The false positive rate was only 1.4%, meaning just 22 real images were mistakenly labeled as fake. This low rate is especially valuable in real-world use because it reduces the likelihood of incorrectly flagging legitimate content. Similarly, the false negative rate stood at 1.8%, with only 15 manipulated images going undetected. While this level is acceptable for most practical scenarios, highly sensitive security applications might still benefit from additional safeguards such as combining multiple models or including manual review for uncertain cases.




Figure 3: Performance Metrics of the Proposed Detection Model

B. Comparative Analysis

To better understand how our model performs in relation to existing research, we carried out a detailed comparison with several well-known baseline approaches using the same test dataset. The methods evaluated included a basic five-layer CNN, the lightweight Meso Net architecture, the deep residual network ResNet-50, and Xception Net, which uses depth wise separable convolutions. Each model was trained under identical conditions, including the same training-testing splits, augmentation techniques, and evaluation procedures, to ensure a fair and unbiased comparison. In addition, hyper parameters for every architecture were carefully tuned through systematic search so that each model operated at its best possible performance level.




Figure 4: Comparative Accuracy Analysis Across Different Detection Methods

Our MobileNetV2-based model outperformed all comparison methods across every evaluation metric, surpassing XceptionNet by 2.87 percentage points, ResNet-50 by 2.67 points, MesoNet by 4.77 points, and the basic CNN baseline by 9.17 points. This notable improvement can be credited to several strengths of the architecture, including its balanced design that maintains both efficiency and generalization, its ability to retain important information through linear bottleneck structures, the use of extensive data augmentation to handle real-world variability, and carefully applied dropout techniques that reduce overfitting during training.

C. System Performance and Efficiency

For practical deployment, strong accuracy must be matched with efficient processing to ensure a smooth user experience. To evaluate this aspect, the system was tested under a variety of real-world conditions, such as single-image predictions, video analysis tasks, batch processing, and scenarios involving multiple simultaneous users. The web platform was also stress-tested over a two-week period with 150 different users completing a total of 2,347 detection requests across a wide range of content types, network environments, and usage behaviors.




Figure 5: System Processing Time Performance Across Different Operations

For single-image analysis, the system required an average of 2.3 seconds from start to finish. This total included approximately 0.4 seconds for file upload of a typical 2MB image, 0.3 seconds for preprocessing, 0.8 seconds for model inference, 0.2 seconds for formatting the results, and about 0.6 seconds for delivering the response, including network delay. Video processing was also efficient, taking about 12.7 seconds to analyze a standard 10-second clip, which corresponds to roughly 1.27 times real-time speed. Within this duration, frame extraction accounted for 3.2 seconds, face detection for 4.5 seconds, CNN prediction for 3.8 seconds, and final aggregation of results for 1.2 seconds. Batch operations further improved efficiency, as processing 10 images together took 18.5 seconds in total—about 1.85 seconds per image—showing a 20% speed improvement compared with sequential execution. Generating detailed reports required only 1.8 seconds, covering compilation of results, layout rendering, and file creation.

Stress testing with as many as 50 simultaneous users demonstrated strong scalability and stable performance. Response time increased gradually from 2.3 seconds with a single user to 3.1 seconds with 50 users, reflecting only a 35% rise despite a fiftyfold increase in workload. This indicates effective handling of resources and request scheduling. During peak usage, GPU

utilization averaged around 75%, suggesting efficient batch processing without overloading the hardware, while memory usage remained steady at about 3.2GB with minimal overhead from memory cleanup processes. Over the testing period, the system maintained 99.7% uptime, with only a few short interruptions for scheduled maintenance and one unexpected restart caused by a database connection issue that was later fixed.

VI. IMPLEMENTATION DETAILS

A. Technology Stack and Infrastructure

The system is built on a modern and reliable technology stack chosen to ensure stability in production, ease of development, and long-term sustainability. The backend is powered by **Django 4.2 running on Python 3.10**, which provides essential web application features such as URL routing, database interaction through an object-relational mapper, template rendering, user authentication, CSRF protection, and an administrative dashboard. Its built-in safeguards against threats like SQL injection and cross-site scripting, along with enforced HTTPS support, strengthen overall application security. In addition, Django's extensive ecosystem of third-party libraries helps streamline the implementation of advanced functionality.

On the frontend, the system uses HTML5 for structured and accessible markup, CSS3 with Flexbox and Grid for responsive layouts across devices, and Bootstrap 5 to maintain consistent design patterns while speeding up development. Interactive behavior is handled through JavaScript with jQuery support for dynamic page updates and asynchronous server communication, reducing the need for full page reloads. The interface follows progressive enhancement principles so that essential features remain accessible even when scripting is unavailable, while enhanced capabilities are provided for modern browsers.

For data storage, the platform relies on MySQL 8.0, which offers strong transactional reliability, optimized query execution, full-text search support, and robust backup and replication tools. The database structure is carefully normalized to reduce redundancy while still maintaining performance through targeted indexing. Persistent connection pooling minimizes overhead and supports high-traffic usage.

Machine learning operations are handled by TensorFlow 2.12 with the Keras API, enabling efficient model training and inference with features such as automatic differentiation, GPU acceleration, distributed computation, and model serialization. OpenCV 4.8 supports image and video processing tasks like format conversion, transformations, face detection, and codec handling, while NumPy 1.24 provides fast numerical operations that underpin the system's data processing.

B. Limitations and Future Scope

Security is designed using a layered approach so that multiple protections work together. User authentication enforces strong password policies requiring sufficient length and character diversity, and real-time strength feedback is provided during registration. Additional safeguards include automatic session expiration after inactivity, temporary account locking after repeated failed login attempts, optional two-factor authentication using time-based codes, and secure password reset procedures verified through email tokens.

User data and uploaded media are protected through several measures. Files are stored using randomized filenames to prevent guessing attacks and can only be accessed by authenticated users with proper permissions. Content is automatically deleted after 30 days unless saved intentionally, and backups are encrypted using strong encryption standards. All network communication is secured through modern TLS encryption with strict transport security enabled, ensuring that data remains protected from interception. API endpoints also enforce request limits to guard against misuse and denial-of-service attempts, helping maintain system stability and reliability.

VII. CONCLUSION AND FUTURE WORK

A. Conclusion

This study successfully designed and deployed a complete deepfake detection system that delivers state-of-the-art performance while also meeting real-world deployment needs. Built on a MobileNetV2 architecture enhanced with a custom classification head, the model achieved 98.47% accuracy in image classification, outperforming well-known baseline approaches such as basic CNN, MesoNet, ResNet-50, and XceptionNet. For video analysis, the system reached an overall accuracy of about 95% by using detailed frame-level evaluation combined with temporal consistency checks, showing strong reliability across varied media types, quality levels, and manipulation methods.

The solution was implemented as a full end-to-end platform within a Django web framework, making it suitable for production use by both individual users and organizations. Performance testing demonstrated high efficiency, with an average image processing time of 2.3 seconds and video analysis running at roughly 1.27 times real-time speed, ensuring smooth interaction for users. Scalability tests confirmed that the platform can handle up to 50 simultaneous users with only gradual performance reduction, and system monitoring recorded 99.7% uptime during extended testing, indicating dependable operation in real deployment environments.

Major contributions of this work include the development of a highly accurate MobileNetV2-based detection model, the design of a sophisticated video analysis pipeline incorporating temporal validation, the creation of a production-ready web platform with comprehensive administrative and user features, and a detailed comparative evaluation proving superior performance over existing methods. The research also provides a thorough analysis of system behavior in terms of accuracy, efficiency, and scalability.

B. Limitations and Future Scope

Although the system performs strongly, some limitations remain. Detection accuracy still depends on the variety and quality of training data, meaning performance could decline when encountering deepfakes produced with new techniques not represented in the dataset. The current approach mainly targets facial manipulations and does not yet address other emerging forms such as full-body synthesis, voice cloning, or text-to-video generation. In addition, resilience against deliberate evasion tactics, such as adversarial perturbations or sophisticated preprocessing, remains an open research challenge.

Future improvements could involve integrating multimodal analysis that combines visual and audio signals to detect inconsistencies, implementing continual learning strategies so the model can adapt to new manipulation methods without forgetting earlier knowledge, and applying adversarial training to strengthen resistance against attacks. Enhancing explainability to provide clearer reasoning behind predictions could also increase user trust. Further optimization techniques—such as model compression and knowledge distillation—may allow deployment on edge devices for real-time applications. Additional directions include blockchain-based verification for tamper-proof authentication records, cross-modal consistency checks across different media formats, and synthetic data generation to expand training diversity and improve generalization.

VIII. ACKNOWLEDGMENTS

We would like to sincerely thank Rao Bahadur Y Mahabaleswarappa Engineering College for offering the infrastructure, technical resources, and institutional backing that supported this research from start to finish. Our heartfelt appreciation goes to our project mentor and the faculty of the Computer Science and Engineering department for their expert advice, thoughtful suggestions, and consistent motivation throughout every stage of the project. We are also grateful to the open-source community—especially the teams behind TensorFlow, Django, and OpenCV—whose well-designed tools and detailed documentation greatly facilitated our work. Lastly, we appreciate the anonymous reviewers for their perceptive feedback, which helped refine and enhance the overall quality of this manuscript.

IX. REFERENCES

- [1] Goodfellow and colleagues (2014) introduced the concept of Generative Adversarial Networks in the *Advances in Neural Information Processing Systems* proceedings, presenting a novel framework for generative modeling.
- [2] Afchar et al. (2018) proposed **MesoNet**, a lightweight neural architecture designed specifically for detecting forged facial videos, presented at the IEEE Workshop on Information Forensics and Security.
- [3] Rossler and collaborators (2019) developed **FaceForensics++**, a large-scale dataset and benchmark for studying manipulated facial imagery, published in the IEEE International Conference on Computer Vision.
- [4] Chollet (2017) authored *Deep Learning with Python*, a comprehensive resource explaining deep learning principles and practical implementation techniques.
- [5] The Django Software Foundation provides official documentation for the Django web framework, offering guidance for development and deployment (accessed January 2025).
- [6] Simonyan and Zisserman (2015) presented a deep convolutional network architecture for large-scale image recognition at the International Conference on Learning Representations.
- [7] Li and co-authors (2020) introduced **Celeb-DF**, a challenging and extensive dataset aimed at improving DeepFake detection methods, presented at the IEEE Conference on Computer Vision and Pattern Recognition.
- [8] Nguyen et al. (2019) explored the use of capsule networks for identifying manipulated images and videos, detailed in the IEEE International Conference on Acoustics, Speech and Signal Processing.
- [9] Güera and Delp (2018) investigated deepfake detection using recurrent neural network models, described in the IEEE International Conference on Advanced Video and Signal Based Surveillance.
- [10] Sabir and colleagues (2019) proposed recurrent convolutional approaches for detecting facial manipulation in videos, presented at the CVPR Workshops.
- [11] Zhao et al. (2021) developed a multi-attentional deepfake detection technique, discussed at the IEEE Conference on Computer Vision and Pattern Recognition.
- [12] Kumar, Bhavsar, and Verma (2022) introduced a metric-learning-based strategy for deepfake identification, presented at the IEEE Winter Conference on Applications of Computer Vision.