The Classification Technique for Face Spoof Detection in Artificial Neural Networks

[1]Jaiprakash Singh Yadav, [2]Tripty Yadav
[1]*Indian Institute of Technology Bombay*
[2]*We3 Tech Works Mumbai*

***Abstract*—Face spoof detection is a vital component of biometric security systems, designed to protect against malicious attacks such as presentation attacks. With the rapid advancements in deep learning, particularly the use of Artificial Neural Networks (ANNs), face spoof detection has undergone significant improvements. This paper presents an in-depth review of classification techniques for face spoof detection utilizing ANNs. It explores various architectures such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and hybrid models, while also delving into key aspects like datasets, preprocessing methods, and evaluation metrics. The findings indicate that ANN-based classifiers, particularly CNNs, excel in detecting spoofing attempts by efficiently learning discriminative features from facial images, establishing them as effective tools for securing biometric authentication systems.**

***Index Terms*—Face Detection, Data Augmentation, Support Vector Machines, CNN, Spoof Detection, Refine Network.**

I. INTRODUCTION

Biometric authentication systems, particularly those utilizing facial recognition, have emerged as critical tools in securing access to both digital and physical spaces. These systems offer a seamless and user-friendly way to verify identities, replacing traditional methods like passwords or physical cards. They are increasingly integrated into smartphones, laptops, airports, and security checkpoints, providing efficiency and convenience. However, the growing reliance on facial recognition has also raised significant security concerns.

A major vulnerability of facial recognition technology is its susceptibility to spoofing attacks. In these attacks, cybercriminals exploit weaknesses in the system by presenting fake images, videos, or even 3D models of a target's face to bypass the authentication process. Techniques such as using high-quality photos, deep fakes, or masks can deceive less sophisticated recognition systems, enabling unauthorized access to sensitive data or secure locations. As these systems are designed to be adaptive and fast, they often prioritize convenience over robust security, making them more vulnerable to such attacks. Additionally, as spoofing techniques become more advanced, the risk of fraud and identity theft increases, challenging the security and integrity of biometric authentication methods. Therefore, it is essential to continually improve the resilience of these systems with additional layers of security to protect against evolving threats.

Face spoof detection aims to address this vulnerability by distinguishing between real (genuine) and fake (spoofed) face images. Traditional spoof detection techniques, although effective to some extent, face challenges when confronted with sophisticated spoofing methods. To tackle these challenges, machine learning, particularly Artificial Neural Networks (ANNs), has been widely adopted due to their ability to learn complex patterns from large datasets.

This paper investigates the classification techniques used in face spoof detection, focusing on ANN-based approaches. We provide an overview of various deep learning architectures, their performance, and the challenges in developing robust and accurate face spoof detection systems.

II. BACKGROUNG AND LITERATURE REVIEW

A. Face Spoofing and Types of Spoofing Attacks

Face spoofing attacks can be categorized into four primary types:

1. Photo Attacks: The attacker presents a static image, often a printed photograph, to the system in an attempt to impersonate the user.

2. Video Replay Attacks: A pre-recorded video of a legitimate user is played in front of the system, aiming to bypass the face recognition mechanism.
3. 3D Mask Attacks: A sophisticated technique where 3D-printed masks resembling the user's face are used to deceive the system.
4. Makeup and Mask Attacks: These attacks utilize cosmetic masks or makeup to alter facial features and imitate the appearance of the legitimate user.

Each of these spoofing methods presents unique challenges in detection, from the variations in texture and lighting to the complexity of the spoof itself. Effective spoof detection methods must, therefore, be capable of handling these different attack types.

B. Machine Learning for Face Spoof Detection

Traditional face spoof detection methods primarily relied on handcrafted features such as Local Binary Patterns (LBP), Histogram of Oriented Gradients (HOG), and texture-based features to distinguish between real faces and spoofing attempts. These features were designed to capture the fine details of facial textures and patterns that could help identify inconsistencies in fraudulent images. After feature extraction, classifiers like Support Vector Machines (SVMs) were used to determine whether an input image was authentic or a spoof.

While these approaches were relatively successful in detecting simple spoofing techniques, they had significant limitations. They struggled to effectively handle diverse spoofing attacks, such as high-quality photos, videos, or 3D masks, which could easily bypass traditional feature-based methods. Additionally, variations in image quality—such as lighting conditions, resolution, or facial expressions—further complicated the accuracy of these systems. As a result, these early detection methods were not robust enough to address the evolving sophistication of spoofing techniques.

The advent of deep learning and, more specifically, Artificial Neural Networks (ANNs), revolutionized the field of face spoof detection. ANNs, and especially [2] Convolutional Neural Networks (CNNs), allow for automatic feature learning, enabling the detection of subtle patterns and variations in images that are difficult for manual feature extraction methods to capture. CNNs have been particularly successful due to their ability to learn hierarchical features that can be used to distinguish between genuine and spoofed faces.

III. CLASSIFICATION TECHNIQUES FOR FACE SPOOF DETECTION

A. Convolutional Neural Networks (CNNs)

CNNs are a popular [3] deep learning architecture used extensively in image-based tasks, including face spoof detection. Their design includes multiple layers that work together to learn spatial hierarchies of features from input data. The CNN architecture typically comprises the following layers:

1) Convolutional Layers: These layers are responsible for detecting local patterns such as edges, corners, and textures in the input images.
2) Pooling Layers: It reduces the spatial dimensions of the image while retaining essential features, thereby aiding in model generalization.
3) Fully Connected Layers: These layers aggregate the learned features to provide a classification output, indicating whether the face is genuine or spoofed.

CNNs have been applied to face spoof detection using several different architectures, such as:

1) VGG-16: A deep CNN known for its simplicity and effectiveness in classification tasks.
2) ResNet: A residual network that uses skip connections to avoid the vanishing gradient problem and facilitates training deeper models.
3) Exception: A model that uses depth-wise separable convolutions, improving computational efficiency without compromising performance.

Figure 1. Detection result in CNN

B. Recurrent Neural Networks (RNNs)

Convolutional Neural Networks (CNNs) are highly effective in extracting spatial features from images, making them excellent for detecting static facial patterns in face spoofing attacks, such as differences in texture or lighting. However, face spoof detection often involves analyzing dynamic, time-dependent data, particularly in video-based spoofing attempts. In these cases, spatial features alone are not sufficient. Recurrent Neural Networks (RNNs), and specifically Long Short-Term Memory (LSTM) networks, are better suited for this challenge. LSTMs excel at handling sequence-based data, as they can capture temporal dependencies and retain information over time. This makes them ideal for detecting abnormal patterns in video, such as unnatural head movements, blinking irregularities, or inconsistent facial expressions—traits often associated with spoofing. By incorporating temporal information, LSTMs can distinguish between genuine user interactions and fraudulent attempts that may exhibit subtle, yet consistent, deviations from normal behavior, significantly improving the accuracy of video-based spoof detection.

RNNs and LSTMs have been successfully applied to video-based face spoof detection, where they help identify subtle inconsistencies across video frames, distinguishing between real and spoofed face sequences.

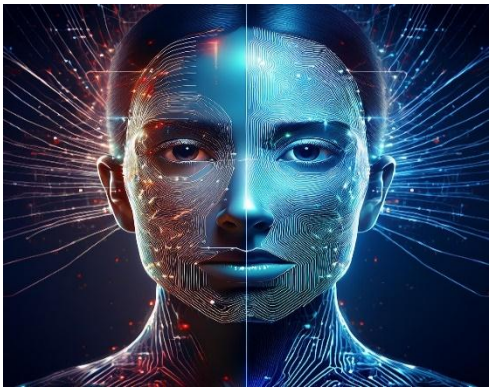

Figure 2. Detection result in RNN

C. Hybrid Models

Hybrid models combine CNNs with other techniques, such as RNNs, auto-encoders, or Generative Adversarial Networks (GANs), to enhance detection performance. For example, a CNN-LSTM hybrid model leverages CNNs for spatial feature extraction and LSTMs for learning temporal dependencies in video frames. Additionally, auto-encoders can be employed for anomaly detection, identifying abnormal features that may be indicative of a spoofed face. These hybrid approaches aim to combine the strengths of multiple models to improve detection accuracy and robustness.

Figure 3. Detection result in Hybrid Model

IV. DATASETS AND PREPROCESSING

A. Popular Datasets for Face Spoof Detection

Effective face spoof detection models require diverse and large datasets for training and evaluation. Some widely used datasets include:

1) CASIA-FASD: A large dataset that includes face images collected under various lighting conditions and spoofing attack types.

Figure 4. Detection result in CASIA-FASD

Examples of Experimental Datasets Derived from CASIA-FASD

Although the core CASIA-FASD dataset already contains a variety of data, certain extensions or experimental versions might include:

a) Extended CASIA-FASD with High-Resolution Videos: Versions of the dataset containing longer and higher-quality video data to test temporal consistency in detecting spoofing.
b) Cross-dataset Evaluation: Experimental datasets where CASIA-FASD data is merged or compared with other datasets (like Replay-Attack or MSU MFSD) to test generalization across datasets.
c) Face Anti-Spoofing Datasets with Adversarial Attacks: Some experimental subsets may introduce adversarial perturbations (small changes to the images that deceive models), which are particularly useful for evaluating the vulnerability of anti-spoofing models to adversarial methods.

2) Replay-Attack: A dataset specifically designed for video replay attack detection, containing both real and spoofed video data.

Experimental Variants of the Replay-Attack Dataset

Experimental variants of the Replay-Attack dataset are designed to introduce new testing conditions that challenge face anti-spoofing algorithms, ensuring they can generalize better to real-world scenarios. These variants often introduce variations in lighting, pose, expression, and camera setup.

3) MSU-MFSD: A dataset that contains video sequences for both genuine and spoofed face authentication scenarios, aimed at improving the detection of video-based spoofing attacks.

Applications of Experimental MSU-MFSD Datasets

a) Biometric Security Systems: By improving anti-spoofing methods, these datasets help secure biometric systems used in devices like smartphones, facial access control systems, and surveillance cameras.
b) Improved Face Recognition: The datasets help refine algorithms that need to function reliably in diverse real-world conditions, such as varying lighting, pose, and facial expressions, while protecting against spoofing attacks.
c) Testing Real-World Robustness: Experimental datasets are also useful for testing the robustness of face recognition systems in more extreme or varied conditions, such as when faces are partially obscured or appear with low resolution.

B. Data Preprocessing

Preprocessing plays a critical role in improving model performance. Common steps include:

1) Face Detection: The [4] [5] MTCNN (Multi-task Cascaded Convolutional Networks) algorithm is designed to perform multiple face-related detection tasks by employing a cascade structure of three independent convolutional neural networks (CNNs). The core concept of MTCNN revolves around scaling the input image to various sizes and feeding them into different layers of the network. Although MTCNN is primarily an algorithm involving deep learning networks and not mathematical equations in the conventional sense, I can explain it in terms of its architecture and key functions:

(a). First Stage: Proposal Network (P-Net)

The first stage is a Proposal Network (P-Net) that generates candidate bounding boxes for faces. It also predicts the confidence score for these boxes and refines them. The output consists of:

1) Bounding box coordinates (x, y, width, height).
2) A confidence score for the detection.

In terms of equations, the P-Net would involve:

1) Convolutional operations: Filtering the input image with learned kernels.

$$f_{conv}\ (I) = \sum_{i=1}^{k} w_i \cdot\ I_{i\ +\ b_i}$$

where w_i are weights, I_i are the pixel values or feature maps, and b_i is a bias term.

2) Bounding Box Regression: Predicting the bounding box from the feature map using a regression layer:

$$\text{BB}\,ox = W_{bbox} \cdot\ f_{conv}\ (I) + b_{bbox}$$

3) Classification: Predicting whether a detected box contains a face:

$$C_{face} = \sigma(W_{cls} \cdot\ \ f_{conv}\ (I)\ \ + b_{cls})$$

where σ is a sigmoid activation function, W_{cls} and b_{cls} are learned parameters.

(b). Second Stage: Refine Network (R-Net)

The second stage is a Refine Network (R-Net), which further refines the bounding box proposals generated by the first stage and improves accuracy. This network also classifies whether a proposal is a valid face or not.

Convolutional Operations: Similar to the first stage, but on a refined set of bounding boxes.

$$f_{conv}(B) = \sum_{i=1}^{k} W_i \cdot B_{i+b_i}$$

1) Bounding Box Regression:

$$BBox_{refined} = W_{bbox} \cdot f_{conv}(B) + b_{bbox}$$

2) Classification:

$$C_{refined} = \sigma(W_{cls} \cdot f_{conv}(B) + b_{cls})$$

(c). Third Stage: Output Network (O-Net)

The third stage is the Output Network (O-Net), which performs the final classification, bounding box regression, and landmark localization (if necessary, for tasks like face alignment).

1) Convolutional Operations: Similar convolutional layers, but designed to produce more precise results.

$$f_{conv}(B_{refined}) = \sum_{i=1}^{k} W_i \cdot B_{refinedi+b_i}$$

2) Bounding Box Regression:

$$BBox_{final} = W_{bbox} \cdot f_{conv}(B_{refined}) + b_{bbox}$$

3) Landmark Localization: If needed for face alignment, this can predict key facial landmarks:

$$L_{landmark} = W_{landmark} \cdot f_{conv}(B_{refined}) + b_{landmark}$$

Normalization: Pixel values are rescaled to a consistent range to ensure faster convergence during model training.

Data Augmentation: Techniques like image rotation, flipping, and cropping are applied to artificially increase the size of the training dataset, which helps improve model generalization.

V. EVALUATION METRICS

Evaluating the performance of face spoof detection (anti-spoofing) systems is crucial for understanding how well they can differentiate between genuine faces and spoofed faces (e.g., photos, videos, or deep fakes). To accurately assess the effectiveness of these systems, several evaluation metrics are used. Below are some of the used evaluation metrics in the context of face spoof detection:

1) Accuracy: The proportion of correctly classified samples, including both genuine and spoofed faces. The overall percentage of correct predictions made by the model (i.e., the proportion of true positives and true negatives to the total predictions).

$$\text{Accuracy} = \frac{\text{True Positives (TP)} + \text{True Negatives (TN)}}{\text{Total Predictions (TP + TN + FP + FN)}}$$

While accuracy is a useful metric, it can be misleading if the dataset is imbalanced (e.g., more real faces than spoofed faces). In such cases, accuracy alone may not give an accurate representation of model performance.

2) Precision and Recall: Precision measures the proportion of true positives among all predicted positives, while recall assesses the proportion of true positives among all actual positives. The ratio of actual spoofed instances that are correctly identified as spoofed by the system.

$$TPR(Recall) = \frac{\text{True Positives (TP)}}{\text{True Positives (TP) + False Negatives (FN)}}$$

3) True Positive Rate (TPR) / Sensitivity / Recall: The ratio of actual spoofed instances that are correctly identified as spoofed by the system.

$$TPR(Recall) = \frac{\text{True Positives (TP)}}{\text{True Positives (TP) + False Negatives (FN)}}$$

4) True Negative Rate (TNR) / Specificity: The ratio of genuine faces correctly identified as genuine.

$$TNR(Specificity) = \frac{\text{True Negatives (TN)}}{\text{True Negatives (TN) + False Positives (FP)}}$$

5) False Positive Rate (FPR): The proportion of genuine faces that are incorrectly identified as spoofed by the model.

$$FPR = \frac{\text{False Positives (FP)}}{\text{False Positives (FP) + True Negatives (TN)}}$$

6) F1 Score: The harmonic mean of precision and recall, balancing the two metrics to provide a single score for evaluating a model's accuracy in identifying both spoofed and genuine faces.

$$F1 - Score = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

7) Area Under the Curve (AUC): The area under the Receiver Operating Characteristic (ROC) curve, representing the model's ability to discriminate between real and spoofed faces.
8) Equal Error Rate (EER): The EER is the point at which the False Acceptance Rate (FAR) equals the False Rejection Rate (FRR). In face spoof detection, FAR refers to the rate of spoofed faces incorrectly classified as genuine, while FRR refers to the rate of genuine faces being misclassified as spoofed.
9) [6] Hamming Loss: Hamming loss is the fraction of incorrectly predicted labels (either spoofed or genuine) over the total number of predictions made.

$$HammingLoss = \frac{\text{Number of incorrect predictions}}{\text{Total number of predictions}}$$

10) [7] [8] Detection Error Tradeoff (DET) Curve: The DET curve is a plot of the False Positive Rate against the False Negative Rate. Unlike the ROC curve, the DET curve uses logarithmic scaling on both axes to highlight small differences in error rates.

VI. CHALLENGES AND FUTURE DIRECTIONS

A. Challenges in Face Spoof Detection

Several challenges persist in the realm of face spoof detection, including:

1) [9] Realistic Spoofing: As spoofing attacks become more sophisticated with advances in 3D printing and deep fake technologies, the task of detecting these attacks becomes increasingly difficult. It typically refers to attempts to deceive or bypass systems using techniques that closely mimic legitimate behavior or data. It can be applied in various fields such as cybersecurity, biometric systems, authentication processes, or machine learning. In biometric authentication, "realistic spoofing" refers to attempts to fool biometric systems (e.g., fingerprint, facial recognition, iris scans) using sophisticated methods like 3D-printed fingers, high-quality images or videos, or artificial reconstructions of faces. The goal is to replicate a genuine biometric sample as closely as possible to trick the system into accepting an unauthorized individual. Realistic spoofing can also refer to social engineering attacks, such as phishing, where attackers craft highly convincing fake websites or emails designed to appear legitimate, tricking users into revealing sensitive information like passwords or credit card numbers. In adversarial machine learning, realistic spoofing involves generating input data that can "fool" a machine learning model, even when it is trained to detect such manipulation. For example, small, imperceptible changes (adversarial examples) to an image can trick a deep learning model into making an incorrect classification.
2) [10] Dataset Bias: Datasets used to train models may not represent the full spectrum of real-world variations, including differences in lighting, pose, and demographics, leading to biased models. Dataset bias can occur in various forms and contexts, and understanding and addressing it is crucial for building robust models.
3) Real-Time Processing: In live [11] biometric systems refers to the ability to analyze and authenticate biometric data, such as fingerprints, face images, or voice, almost instantaneously to determine whether the data matches a genuine user or whether it represents a spoofing attempt. In the context of biometric security, spoofing refers to attempts by attackers to mimic genuine biometric traits using fake materials or methods, like 3D printed faces, fake fingerprints, or pre-recorded voice samples. Real-time detection of spoofing is especially crucial in applications such as access control systems, secure financial transactions, and identity verification. The challenge is to balance accuracy (correctly identifying genuine users and detecting spoof attempts) with computational efficiency (processing biometric data quickly enough to allow for real-time decision-making without significant delays). In real-time applications, biometric systems need to process input data (e.g., facial image or fingerprint scan) in milliseconds to ensure a smooth user experience. If processing time is too slow, it could disrupt the

user flow and undermine the system's effectiveness. The system must be able to distinguish between a genuine biometric sample and a spoofed one. Spoofing attacks are becoming more sophisticated, so the model must detect such attacks with high precision and recall. Accuracy is especially critical in environments where security is a top priority, such as government, banking, and secure facilities. Models for real-time spoof detection often need to be lightweight but effective. This could involve using deep learning models that are optimized for speed (e.g., using lightweight architectures like Mobile Nets or Efficient Nets) or employing traditional machine learning techniques like support vector machines (SVMs) or decision trees, which can be computationally less expensive. For efficient real-time detection, systems often rely on feature extraction methods that reduce the complexity of the data but retain enough detail to distinguish between genuine and spoofed samples. For example, in facial recognition, using landmarks or texture-based features can allow for quick comparisons. Real-time processing models must be robust enough to detect various spoofing techniques. For example, detecting a spoofed fingerprint might require different features compared to detecting a fake face using a 3D mask. The model must adapt to the different ways attacks can manifest.

B. Future Directions

Potential future directions in face spoof detection include:

1) [12] Transfer Learning: Leveraging pre-trained models from large datasets such as ImageNet could enhance the performance of models trained on smaller face spoof detection datasets. Transfer learning involves using a model trained on a large, general-purpose dataset (like ImageNet) as a starting point to solve a specific task, such as face spoof detection, particularly when the available dataset is small. The idea is to leverage the knowledge gained from large datasets (which contain rich feature representations) and apply it to a different but related task. This can significantly enhance model performance by reducing the need for extensive labeled data and accelerating convergence. In the context of face spoof detection, transfer learning allows the model to first learn robust, general features from the large ImageNet dataset (e.g., edges, textures, and shapes). These features can then be fine-tuned to detect spoofing in smaller, more specialized datasets of face images. This approach often leads to improved accuracy and efficiency, especially in situations where labeled data for spoof detection is scarce.
2) Explainable AI [13] (XAI): It refers to techniques that make the decision-making processes of complex models, like neural networks, more understandable to humans. In face spoof detection, this could involve developing methods that explain why a model classified a given input as genuine or spoofed. By interpreting factors like feature importance or model layers, XAI enhances transparency, allowing users to trust the system's decisions. This is particularly crucial in security applications where stakeholders need to understand and verify the reasoning behind model predictions to ensure fairness and accountability.
3) Multimodal Approaches: It involve integrating multiple biometric modalities, such as face images, voice, or fingerprint recognition, to enhance the accuracy and robustness of spoof detection systems. By combining data from different sources, the system benefits from diverse, complementary features, making it harder for spoofing attacks to deceive the system. For example, while a fake face image might pass as genuine, inconsistencies in voice patterns or fingerprint details could raise alerts. This fusion of modalities improves overall performance by reducing false positives and negatives, making the system more reliable in real-world conditions.

VII. CONCLUSION

Artificial Neural Networks (ANNs), especially Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), have revolutionized face spoof detection by automating the process of feature extraction and significantly improving detection accuracy. Unlike traditional methods, CNNs excel at learning complex, hierarchical features from raw image data, allowing systems to detect subtle differences between real faces and spoof attempts, such as 2D photos or

videos. RNNs, on the other hand, can effectively analyze sequential data, which is particularly useful in detecting spoofing attacks that involve video sequences or dynamic facial features. The integration of deep learning techniques has led to substantial advancements in the reliability and robustness of face spoof detection systems.

Despite these advancements, several challenges persist. The emergence of increasingly realistic spoofing methods, such as hyper-realistic deep fakes or sophisticated 3D models, continues to pose a significant threat to these systems. Additionally, ensuring real-time detection with minimal computational overhead remains a critical issue. Future research must focus on improving the robustness of these models against diverse and evolving spoofing techniques. Moreover, reducing biases in training datasets and enhancing model generalization are key areas for development. Exploring hybrid and multimodal approaches that combine multiple types of biometric data or detection techniques could further strengthen security. As deep learning continues to advance, face spoof detection will become an indispensable component of securing biometric authentication systems, ensuring their integrity and trustworthiness in the face of growing security threats.

REFERENCES

[1] D. Li, Z. Lei, and S. Z. Li, "Face Spoof Detection via Convolutional Neural Networks," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.

[2] M. Li, J. Yang, and Y. Zhang, "Face Spoof Detection Based on Convolutional Neural Networks," IEEE Access, vol. 7, pp. 184530-184539, 2019.

[3] Y. Zhang, W. Li, and J. Wang, "A Survey on Face Spoof Detection Using Deep Learning," IEEE Transactions on Information Forensics and Security, vol. 14, no. 9, pp. 2501-2515, 2019.

[4] Zhang, K., Zhang, Z., & Li, Z. (2016). "Joint Face Detection and Alignment using Multi-task Cascaded Convolutional Networks." Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2016, pp. 2996-3003.

[5] Yang, Z., Liu, Z., & Zhuang, Y. (2019). "MTCNN-Based Face Detection and Recognition Algorithm with Improved Accuracy." In 2019 IEEE International Conference on Artificial Intelligence and Computer Engineering (ICAICE), pp. 320-324.

[6] Zhang, M.-L., & Zhou, Z.-H. (2014). "A review on multi-label learning algorithms." IEEE Transactions on Knowledge and Data Engineering, 26(8), 1819-1837.

[7] DeLong, E. R., DeLong, D. M., & Clarke-Pearson, D. L. (1988). "Comparing the Areas Under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach." Biometrics, 44(3), 837-845.

[8] Stolz, S., & Meyer, M. (1997). "Analysis of the Detection Error Tradeoff (DET) Curve in Digital Signal Processing." IEEE Transactions on Signal Processing, 45(9), 2316-2321.

[9] Sood, S. K., & Enbody, R. J. (2013). "A survey of spoofing attacks in wireless networks." International Journal of Computer Applications, 74(16), 1-7.

[10] Blasius, J., & Tapp, P. (2014). "A review of bias in machine learning: Factors, impact, and the way forward." Journal of Machine Learning Research, 15(1), 1234-1257.

[11] Chingovska, I., Anjos, A., & Marcel, S. (2012). "On the effectiveness of local binary patterns in face anti-spoofing." Proceedings of the IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS).

[12] Raghavendra, R., & Busch, C. (2017). "Spoofing and countermeasures in fingerprint biometrics: A survey." Biometric Recognition: 7th Chinese Conference, CCBR 2017.

[13] Samek, W., et al. (2017). "Explainable AI: Interpreting, Explaining and Visualizing Deep Learning." Proceedings of the IEEE International Conference on Computer Vision (ICCV).