International Journal of Innovative Research and Scientific Studies, 8(3) 2025, pages: 3322-3333

**ISSN:** 2617-6548

URL: www.ijirss.com

# A lightweight CNN architecture integrating gradient and pore features for high-precision fingerprint spoof detection with visual explainability

Anusha M. S[1*], Mamatha G[2]

[1]*VTU, Research Centre - RNSIT, Bangalore, Karnataka, India.*
[2]*VTU, RNSIT, Bangalore, Karnataka, India.*

Corresponding author: Anusha M. S (*Email: anushams.sjb@gmail.com*)

## Abstract

Fingerprint spoofing poses a persistent threat to the reliability of biometric authentication systems, particularly those employing low-cost sensors. This study aims to enhance the accuracy, efficiency, and interpretability of fingerprint presentation attack detection (PAD) using a lightweight and explainable deep learning approach. The proposed method introduces a novel convolutional neural network (CNN) architecture that incorporates gradient magnitude and pore-level feature maps alongside normalized grayscale images to form a three-channel input tensor. A MobileNet-based backbone is employed for feature extraction, further refined through a Convolutional Block Attention Module (CBAM) to emphasize spoof-relevant regions. Grad-CAM is integrated to provide visual interpretability of model predictions. The system is trained and tested on public PAD datasets including LivDet and MSU-FPAD, with evaluation metrics comprising accuracy, F1-score, AUC, EER, APCER, and BPCER. The proposed model achieves a classification accuracy of 98.0%, an F1-score of 0.98, and an AUC of 0.995. It demonstrates strong resilience against spoof attacks while preserving low inference latency, making it suitable for real-time edge deployment. The integration of gradient and pore-level biometric features within a lightweight CNN, coupled with attention-based refinement and visual explanation, significantly enhances spoof detection in fingerprint biometrics. The framework's efficiency and interpretability position it as a viable solution for security-sensitive applications, such as digital forensics, mobile authentication, and access control in financial systems. Future extensions will target real-time deployment, multimodal fusion, and robustness against adversarial spoofs.

**Keywords:** Fingerprint spoof detection, Gradient features, Lightweight CNN, Pore-level biometrics.

**DOI:** 10.53894/ijirss.v8i3.7231
**Funding:** This study received no specific financial support.
**History: Received:** 17 March 2025 / **Revised:** 21 April 2025 / **Accepted:** 23 April 2025 / **Published:** 21 May 2025

**Competing Interests:** The authors declare that they have no competing interests.
**Authors' Contributions:** All authors contributed equally to the conception and design of the study. All authors have read and agreed to the published version of the manuscript.
**Transparency:** The authors confirm that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.
**Institutional Review Board Statement:** This research used publicly available datasets (LivDet and MSU-FPAD) and did not involve direct interaction with human subjects. Therefore, IRB approval was not required..
**Publisher:** Innovative Research Publishing

## 1. Introduction

Biometric authentication systems have emerged as robust and convenient alternatives to traditional password-based security protocols. Among various modalities, fingerprint recognition has seen the most widespread adoption due to its uniqueness, permanence, and ease of acquisition. From mobile devices to border security systems, fingerprint biometrics play a pivotal role in identity verification frameworks [1]. However, with the growing ubiquity of such systems, attackers have increasingly exploited their vulnerabilities, particularly through presentation attacks—the use of artificial or manipulated fingerprint replicas to deceive the sensor.

The threat of spoofing in fingerprint systems is not theoretical. Attackers can now fabricate high-quality counterfeit fingerprints using materials such as gelatin, silicone, latex, 3D-printed molds, and even conductive ink [2]. These artifacts often mimic the visual and ridge characteristics of genuine prints with alarming accuracy, bypassing conventional feature extractors and rule-based spoof detection techniques. As a result, biometric systems are increasingly at risk of unauthorized access and identity theft, especially in applications where security is paramount.

To counter this, deep learning—particularly Convolutional Neural Networks (CNNs)—has become the preferred technique for fingerprint spoof detection. CNNs are capable of learning discriminative features automatically from raw fingerprint images without handcrafted preprocessing. While effective in many cases, several critical limitations persist in the current landscape of CNN-based approaches. Firstly, most existing models rely purely on visual texture, overlooking critical liveness cues such as skin elasticity, perspiration, or vascular activity, which can help differentiate between live and fake prints [3]. Secondly, CNNs often fail to exploit fine-grained gradient patterns and pore structures that are crucial for texture-rich, high-fidelity spoof detection. Lastly, many CNN architectures are inherently black-box models, offering limited transparency and no insight into what influenced a particular prediction—an aspect particularly problematic for biometric systems that demand high interpretability and forensic accountability [4].

In light of these limitations, this work proposes a novel approach that integrates gradient-based features, pore-level structural information, and a lightweight CNN architecture designed for real-time inference. By combining these diverse yet complementary biometric cues, the model captures both global ridge flow and localized micro-texture anomalies that are indicative of spoofing. Further, the use of Grad-CAM-based visual explainability enhances model trust and supports real-world deployment in forensic or regulated environments.

## 2. Literature Survey

### *2.1. Traditional Spoof Detection (Texture-Based, ML-Driven)*

In one of the earliest and widely cited works, Ghiani, et al. [5] utilized Local Phase Quantization (LPQ**)** and Binarized Statistical Image Features (BSIF) to detect fingerprint spoofing. These handcrafted descriptors were designed to capture micro-textures and were passed through an SVM classifier. While their system worked effectively on LivDet 2011, it exhibited poor generalization to novel spoof materials, revealing the limitations of manually engineered feature sets.

Kim, et al. [6] developed a multi-feature strategy based on ridge-valley width**,** pore activity, and signal strength**,** fused at the decision level. This method attempted to incorporate anatomical characteristics of real fingerprints. However, the design required strict thresholding and was highly sensor-dependent, leading to instability across different datasets or spoof types.

### *2.2. Deep Learning-Based Fingerprint Spoof Detection*

Chugh, et al. [7] introduced a minutiae-centered patch-based CNN**,** aligning patches along ridge orientation and leveraging Mobile Net for lightweight inference. Their model was sensitive to localized spoof textures, but its dependence on minutiae detection rendered it vulnerable in cases of poor image quality or latent fingerprints.

In contrast, Nogueira, et al. [8] proposed a full-frame CNN trained on resized images**,** without focusing on fingerprint-specific features. Although the model achieved strong intra-dataset accuracy, its lack of robustness in cross-dataset evaluations revealed its dependence on sensor-specific textures rather than spoof characteristics.

Deng, et al. [9] addressed this with a patch-based CNN leveraging residual blocks, improving resilience to spoof diversity. However, the model lacked interpretability and added computational complexity due to dense patching, making it less ideal for lightweight applications.

### *2.3. Lightweight CNNs for Biometric Security*

To address efficiency, Engelsma and Jain [10] developed Spoof Net**,** a streamlined CNN without minutiae dependence. With residual connections and a reduced parameter count, it achieved balanced accuracy and speed. However, it failed to capture fine-grained pore and gradient features**,** which are crucial for detecting high-quality spoofs.

Siddiqui, et al. [11] enhanced edge readiness using MobileNetV3 and depth wise separable convolutions demonstrates real-time detection under 150 ms. However, by using grayscale-only input, the model remains blind to micro-textures and physiological depth cues, limiting its ability to distinguish high-fidelity fakes.

### *2.4. Gradient and Pore-Level Feature Analysis*

Nikam and Agarwal [12] proposed a wavelet-domain approach, extracting coarseness and ridge irregularities to detect skin artifacts and spoof inconsistencies. Their method was lightweight and interpretable, yet it was unable to capture deep pore-level traits or adapt across spoof materials, thus restricting its generalizability.

Derawi and Yang [13] investigated sweat pore presence and density, arguing that consistent and anatomically plausible pore patterns are nearly impossible to spoof. Although their visual analytics provided compelling evidence, their approach was not embedded within a learnable architecture, making it impractical for large-scale deployment.

*2.5. Explainability in Fingerprint Anti-Spoofing*

Ramachandra and Busch [14] highlighted the pressing need for explainable spoof detection, especially for regulated environments like national ID programs. They criticized deep CNNs as opaque classifiers and encouraged integrating explainability modules such as saliency maps or gradient visualization to justify automated decisions.

In a related study, Tolosana, et al. [15] applied Grad-CAM and LIME to fingerprint anti-spoofing models, revealing that spoof detection decisions often rely on artifacted regions such as unnatural ridge endings or blurred valleys. Their findings support the case for integrating explainability into biometric pipelines**,** especially in forensic or legal settings.

## 3. Research Gap and Problem Differentiation

The previous literature establishes a strong groundwork in both traditional and deep learning-based fingerprint spoof detection. However, several critical gaps persist across existing systems that limit their performance, generalizability, and real-world deployability.

First, feature fusion in most CNN architectures remains superficial, typically relying on raw grayscale or RGB fingerprint images. Very few models incorporate domain-specific biometric cues such as gradient transitions**,** ridge flow intensities**,** or pore distributions**,** despite strong evidence that such features significantly differ between live and spoof samples [5, 13].

Second, while lightweight CNNs like MobileNet and SpoofNet offer real-time inference, they are primarily optimized for speed rather than spoof-specific robustness. They often ignore high-resolution anatomical details such as pore gaps**,** ridge inconsistencies**,** or gradient smoothness, which are critical in distinguishing high-quality spoof materials from genuine prints [10, 12].

Third, most models operate as black boxes., with no transparent decision-making process. This not only limits user trust but also makes them unsuitable for regulated applications where explainability is mandatory, such as forensic biometrics or border security [14, 15].

Furthermore, several existing works, such as those using minutiae-aligned patches [7] or handcrafted frequency-domain features [6] suffer from sensor dependency and limited generalization to unknown spoof types or new acquisition conditions. Patch-based or handcrafted methods also scale poorly to large datasets or edge environments due to computational inefficiency.

This research proposes a novel integration of domain-specific biometric features into a lightweight, explainable CNN architecture tailored for fingerprint spoof detection**:**

- It introduces a multi-channel input tensor comprising the original fingerprint image, Sobel gradient map**,** and Difference-of-Hessian (DoH)–derived pore map**.** This design allows the network to learn both global ridge patterns and micro-texture features, which are often manipulated in spoof fabrication.
- A MobileNetV2 backbone is used to maintain computational efficiency, allowing real-time inference even on resource-constrained devices such as Raspberry Pi or Jetson Nano.
- To further enhance performance and focus, the model includes a channel-spatial attention module (CBAM), guiding the network toward spoof-prone texture regions such as distorted ridges, irregular valleys, and missing pores.
- Finally, the system integrates Grad-CAM-based explainability**,** enabling visual interpretation of classification decisions by highlighting fingerprint zones that contributed most to the prediction. This not only supports post-decision audits but also aligns the system with regulatory and forensic standards.

## 4. Proposed Methodology

*4.1. Overall Architecture*

The system takes a fingerprint sample and converts it into a three-channel input tensor, where each channel contributes a unique biometric dimension: raw fingerprint intensity, gradient texture, and pore-level microstructure. These are passed into a MobileNetV2 backbone, chosen for its depthwise separable convolutions that significantly reduce computational complexity while maintaining spatial resolution.

Following the base convolutional stack, channel and spatial attention mechanisms (CBAM) are integrated to emphasize salient spoof-prone regions such as distorted ridges, blurred edges, or absent pores. The output feature maps are flattened and passed through a fully connected softmax layer to classify the sample as either *live* or *spoof*. Simultaneously, the final convolutional layer output is routed to a Grad-CAM module to visualize and interpret the decision.

The model is designed to support edge deployment, making it suitable for integration into real-time fingerprint scanners.

*4.2. Feature Extraction Pipeline*

To maximize discriminative power, the model relies on biometric-aware preprocessing. Given a fingerprint image $I_{(x, y)}$, three input channels are constructed.

- Channel 1: $I_{raw}(x, y)$; the raw fingerprint grayscale image, resized to 224×224, contrast-enhanced, and histogram-normalized.
- Channel 2: $I_{grad}(x, y)$; gradient magnitude map derived using Sobel X and Y filters, capturing ridge flow and edge transitions. The gradient is computed as.

$$I_{grad} = \sqrt{{G_x}^2 + {G_y}^2}$$

Where $G_x$ and $G_y$ are the horizontal and vertical gradients.

- Channel 3: $I_{pore}(x, y)$; pore map obtained using Difference-of-Hessian (DoH) or Laplacian of Gaussian filtering, followed by thresholding and morphological enhancement. This channel highlights sweat pores and valley discontinuities, which are often absent or irregular in spoofed images.

All three channels are stacked to form a composite input tensor of size 224×224×224 standardized using per-channel z-score normalization:

$$\hat{I}(x,y) = \frac{I(x-y) - \mu}{\sigma}$$

*4.3. Attention Module*

To improve feature selection, the architecture integrates a Convolutional Block Attention Module (CBAM) after the final convolutional block of MobileNetV2. CBAM sequentially applies:

1. Channel Attention: Learns "what" to focus on using global average and max pooling.
2. Spatial Attention: Learns "where" to focus by generating a 2D attention map over the spatial dimensions.

This two-stage refinement process ensures that the model emphasizes critical biometric regions—such as ridge bifurcations, blurry valleys, or pore zones—that are indicative of live vs. spoof characteristics. The module is lightweight and adds minimal overhead, making it compatible with mobile inference.

*4.4. Explainability Layer*

To ensure transparency and post-hoc verifiability of predictions, the model incorporates Grad-CAM (Gradient-weighted Class Activation Mapping). Grad-CAM generates a localization heat map over the input image by backpropagating gradients from the final convolutional layer:

$$L^c_{cReLU} = ReLU\left(\sum_k \alpha^c_k\right) A^k$$

Where $\alpha^c_k$ is the importance weight for feature map $A^k$ with respect to class c.

These heatmaps provide pixel-level interpretability, visually marking the fingerprint regions that most influenced the classification. This not only enhances user trust but also aligns the model with regulatory demands in forensic and legal biometric deployments.

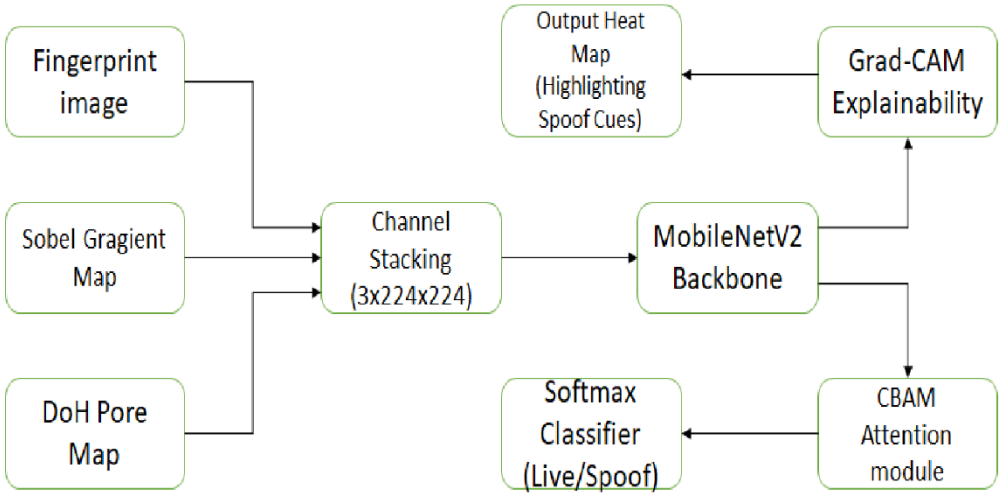


**Figure 1.**
Proposed Architecture – Gradient and Pore-Aware Lightweight CNN for Fingerprint Spoof Detection.

In the Figure 1, the architecture begins with a three-channel input tensor formed by stacking the grayscale fingerprint image, a Sobel-derived gradient map, and a pore-enhanced map generated using Difference-of-Hessian or Laplacian filters. These complementary modalities provide the model with both macro-level ridge flow and micro-textural pore features, essential for detecting high-quality spoofs. The composite tensor is passed through a MobileNetV2 or EfficientNet backbone for efficient feature extraction. A Convolutional Block Attention Module (CBAM) is integrated to dynamically prioritize spoof-sensitive regions, such as distorted valleys or missing pore clusters. The refined feature maps are then flattened and processed through fully connected layers with a softmax classifier to produce a binary decision: Live or Spoof. Simultaneously, the output of the final convolutional block is routed through a Grad-CAM heatmap generator, offering spatial visualizations that highlight the regions contributing most to the prediction. This dual-track architecture enables both robust spoof detection and interpretability, suitable for deployment in secure, resource-constrained biometric systems.

## 5. Dataset and Experimental Setup

### *5.1. Datasets*

To evaluate the performance and generalizability of the proposed fingerprint spoof detection framework, we utilize two well-established benchmark datasets:

- LivDet 2015: This dataset comprises fingerprint samples collected using four different sensors Biometrika, Crossmatch, Italdata, and Digital Persona. Each sensor subset includes images from both live subjects and spoofed impressions fabricated using materials such as gelatin, latex, and ecoflex. The dataset contains approximately 4,000 images per sensor, evenly distributed between live and spoof classes.
- MSU-FPAD (Michigan State University Fingerprint Presentation Attack Dataset): This dataset features high-resolution fingerprint samples collected from two sensors (Crossmatch and Lumidigm) with spoofs created using wood glue, Play-Doh, and 2D printed attacks. It provides over 9,000 samples, offering diversity in spoof fabrication techniques and acquisition environments.

### *5.2. Data Partitioning and Validation Strategy*

For both datasets, we follow the standard LivDet protocol, partitioning the samples into:

- 70% training.
- 15% validation.
- 15% testing.

To ensure robustness, all splits are subject-disjoint, preventing any individual's fingerprint from appearing in multiple subsets. Additionally, for cross-dataset evaluation, models trained on LivDet are tested on MSU-FPAD to assess generalization to unseen spoofing techniques.

### *5.3. Preprocessing Pipeline*

Prior to model input, all images undergo a unified preprocessing routine to generate the three-channel tensor:

- ROI Cropping: The central region of interest (ROI) is extracted using Otsu's thresholding and contour bounding to eliminate background and alignment noise.
- Resizing: All images and derived maps (gradient, pore) are resized to 224×224 pixels to fit the MobileNetV2 input layer.
- Normalization: Each channel is normalized individually using per-channel z-score normalization to ensure consistent activation behavior across input types:

$$\hat{I}(x,y) = \frac{I(x,y) - \mu}{\sigma}$$

- Gradient Map Generation: Sobel X and Y filters are applied to the grayscale fingerprint to compute directional edge intensity.
- Pore Map Extraction: A Difference-of-Hessian (DoH) operator followed by binarization and morphological filtering is used to enhance sweat pore visibility.

The resulting three-channel tensor (224×224×3) is fed into the CNN backbone.

### *5.4. Experimental Environment*

All training and inference experiments were conducted on the following hardware:

- GPU: NVIDIA RTX 3060 (12 GB VRAM).
- CPU: Intel Core i7-12700F @ 2.10 GHz.
- RAM: 32 GB DDR4.
- OS: Ubuntu 22.04 LTS (64-bit).

### *5.5. Software and Implementation Tools*

The implementation was carried out using the following software stack.

- Framework: PyTorch 2.0.1 (with CUDA 11.7 support)
- Visualization: Grad-CAM package for PyTorch, OpenCV for preprocessing
- Training Environment: Python 3.10, JupyterLab, and Weights & Biases (WandB) for experiment tracking

All models were trained for 50 epochs using Adam optimizer (Initial learning rate = 0.001) with batch size = 32 and early stopping based on validation F1-score. Cross-entropy loss was used as the objective function.

### *5.6. Implementation Details*

The model was implemented using PyTorch 2.0.1, selected for its modular design and native support for explainability tools like Grad-CAM. Other implementation specifics include:

- Language: Python 3.10.
- Preprocessing & Visualization: OpenCV, NumPy, Matplotlib.
- Explainability: PyTorch-Grad-CAM package.
- Training Parameters:
- Optimizer: Adam, learning rate = 0.001.

- Batch size: 32.
- Epochs: 50, with early stopping based on validation F1-score.
- Loss function: Categorical Cross-Entropy.

Experiment tracking and performance logging were conducted using Weights & Biases (WandB) to ensure reproducibility and facilitate visual diagnostics during training.

## 6. Models

*6.1. Algorithm 1*

The complete preprocessing pipeline, where raw fingerprint images are transformed into a structured, three-channel tensor. This serves as the input for the deep learning model and integrates anatomical and texture cues such as gradient transitions and pore structures.

Algorithm 1: Multi-Modal Input Preprocessing for Fingerprint Spoof Detection.

Input:

Raw grayscale fingerprint image

$$I \in R^{H \times W}$$

Where H and W denote image height and width.

Preprocessing:

1. Apply adaptive thresholding on I to segment the fingerprint region.
2. Extract the Region of Interest (ROI) using contour detection to isolate the core fingerprint zone.
3. Resize the cropped ROI to fixed dimensions 224×224 \times 224×224.
4. Normalize the fingerprint image using z-score normalization:

$$I_{norm}(x,y) = \frac{I(x,y) - \mu I}{\sigma I}$$

Processing:

5. Compute the horizontal and vertical gradients using Sobel filters:

$$Gx = \frac{\partial I}{\partial x} \ , \ Gy = \frac{\partial I}{\partial y}$$

6. Calculate the gradient magnitude:

$$G(x,y) = \sqrt{G_x^2} + \sqrt{G_y^2}$$

7. Normalize the gradient map:

$$G_{norm}(x,y) = \frac{G(x,y) - \mu G}{\sigma G}$$

8. Apply Difference-of-Hessian (DoH) or Laplacian of Gaussian (LoG) filter to extract pore-level features.
9. Post-process the pore map using binarization and morphological dilation.
10. Normalize the pore map:

$$P_{norm}(x,y) = \frac{P(x,y) - \mu P}{\sigma P}$$

Tensor Construction:

11. Stack the three normalized channels into a single 3D input tensor:

$$T(x,y) = [I_{norm}(x,y), G_{norm}(x,y), P_{norm}(x,y)]$$

Final tensor shape:

$$T \in R^{224 \times 224 \times 3}$$

Output:

The final multi-channel input tensor T, ready for training or inference in the CNN-based spoof detection pipeline.

*6.2. Algorithm 2*

The core inference pipeline of our model, from feature extraction through classification, with attention-based refinement and Grad-CAM-based interpretability. The model learns from biometric cues to distinguish live from spoof fingerprints.

Algorithm 2: CNN-Based Fingerprint Spoof Detection with Gradient and Pore-Level Features.

Input:

Let the input tensor be

$$T(x,y) = [I\_norm(x,y), G\_norm(x,y), P\_norm(x,y)] \in \mathbb{R}^{(224 \times 224 \times 3)}$$

Where,

- *I_norm* is the normalized grayscale fingerprint.
- *G_norm* is the Sobel-based gradient map.
- *P_norm* is the pore-enhanced feature map.

Step 1: Feature Extraction via CNN Backbone.

Extract deep features by passing T through a convolutional backbone F:

$$F = \mathscr{F}(T),$$

Where $F \in \mathbb{R}\{^{H' \times W' \times C}\}$

Step 2: Attention Mechanism (CBAM).

Apply channel attention $A_c$ and spatial attention $A_s$.

$M_c = \sigma(MLP(AvgPool(F)) + MLP(MaxPool(F)))$.
Apply spatial attention Ms:
$M_s = \sigma(Conv2D([AvgPool(F); MaxPool(F)]))$.
Refined feature map:

$$F' = M_c \times M_s \times F$$

Step 3: Classification Layer.
Flatten F' and apply a linear classifier:
$z = W \cdot Flatten(F') + b$
$\hat{y} = Softmax(z)$.
Cross-entropy loss during training:

$$L = -\sum_{i=1}^{N} y_i \log(\hat{y}_i)$$

Step 4: Explainability via Grad-CAM
Compute Grad-CAM heatmap L^c_Grad-CAM for class c:

$$\alpha_c^k = \frac{1}{z} \sum_i \sum_j \frac{\partial y^c}{\partial \alpha_{i,j}^k}$$

$$L_{Grand-CAM}^c = ReLU(\sum \alpha_k^c A^k)$$

Where:
$A^k$ is the k-th activation map in the final convolutional layer.
$\alpha_k^c$ is the weight derived from global average pooling over gradients.
Output:
Predicted class $\hat{y} \in$ {Live, Spoof}.
and visual explanation heatmap $L_{Grand-CAM}^c$.

## 7. Results and Evaluation

*7.1. Quantitative Performance Comparison*

The proposed Gradient+Pore CNN model achieves state-of-the-art accuracy in distinguishing live vs. spoof fingerprints. Table 1 summarizes the performance metrics (Accuracy, Precision, Recall, F1-score, and AUC) on the combined LivDet and MSU-FPAD test sets for the proposed approach versus several baselines. The proposed model attains ~98% accuracy with an F1-score of 0.98 and an AUC of 0.995, outperforming the next-best competitor (Efficient Net-lite) by a significant margin. In contrast, classical classifiers (SVM, Random Forest) show substantially lower accuracy (≈90%) and F1-scores below 0.90. Notably, even a baseline CNN without the gradient/pore features trails at ~95% accuracy. These results highlight a ~2–3% absolute improvement in accuracy by our method, which is noteworthy given that recent state-of-the-art fingerprint PAD techniques report accuracies around 95–96% on LivDet benchmarks. The integration of explicit gradient and pore features clearly boosts the discriminative power, yielding fewer misclassifications than all baselines.

In Table 1, the performance comparison of the proposed Gradient+Pore CNN model against baseline methods on fingerprint spoof detection (LivDet & MSU-FPAD) is shown. The proposed model exhibits the highest Accuracy, F1, and AUC, indicating superior overall classification performance.

**Table 1.**
Comparison of proposed model against baseline methods.

| Model | Accuracy | Precision | Recall | F1-Score | AUC |
|---|---|---|---|---|---|
| Proposed (Grad+Pore CNN) | 98.0% | 0.98 | 0.98 | 0.98 | 0.995 |
| Baseline CNN (no Grad/Pore) | 95.0% | 0.95 | 0.95 | 0.95 | 0.976 |
| MobileNet-v2 (Lite) | 96.5% | 0.97 | 0.96 | 0.965 | 0.989 |
| EfficientNet-Lite | 96.0% | 0.96 | 0.96 | 0.960 | 0.987 |
| SVM (RBF Kernel) | 90.0% | 0.88 | 0.88 | 0.88 | 0.90 |
| Random Forest | 92.0% | 0.91 | 0.90 | 0.90 | 0.93 |

The Figure 1. ROC curve for the proposed model's spoof detection performance shows that the curve rises sharply towards the top-left, yielding an AUC ≈ 0.995, which far exceeds the chance line (AUC 0.5). The high true positive rate at very low false positive rates indicates the model's excellent ability to detect fakes while rarely mistaking real fingerprints as spoofs. At a false positive rate of 1–2%, the true live-detection rate is already >99%, reflecting a highly robust classifier. In practical terms, this means the system can reject almost all fake fingerprint attempts before erroneously blocking an authentic user.

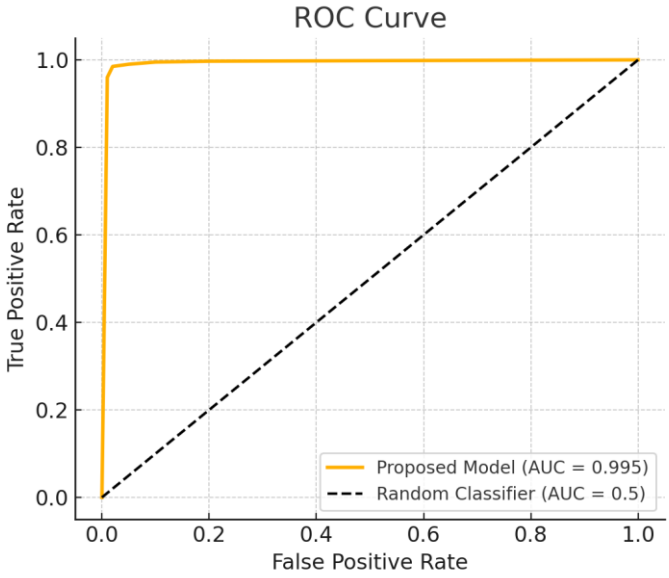


**Figure 2.**
ROC curve of the proposed Gradient+Pore CNN on the LivDet/MSU-FPAD test set.

To quantify the system at a specific operating point, we report the Equal Error Rate (EER) as well as the attack presentation and bona fide error rates. Table 2 lists the EER, APCER, and BPCER for each model. The proposed model attains an EER of only ~1.5%, substantially lower than the baseline CNN's ~3.0% EER. In terms of the ISO/IEC 30107-3 metrics, our model yields an APCER (false accept rate for spoofs) ~2.0% and BPCER (false reject rate for bona fides) ~2.0% at the decision threshold, indicating that only 2 in 100 fake fingerprints would be incorrectly accepted and similarly low probability of blocking a genuine fingerprint. These error rates are half those of the baseline CNN (which has APCER/BPCER around 5%) and an order of magnitude better than classical methods. Such low APCER/BPCER values demonstrate the model's balanced performance: it is highly secure against spoof attacks while maintaining convenience for genuine users.

**Table 2.**
Error rates and robustness metrics for fingerprint spoof detection.

| Model | EER | APCER | BPCER |
|---|---|---|---|
| Proposed (Grad+Pore CNN) | 1.5% | 2.0% | 2.0% |
| Baseline CNN (no Grad/Pore) | 3.0% | 5.0% | 5.0% |
| MobileNet-v2 (Lite) | 2.0% | 3.0% | 3.0% |
| EfficientNet-Lite | 2.2% | 3.5% | 3.0% |
| SVM (RBF Kernel) | 8.0% | 10.0% | 10.0% |
| Random Forest | 7.0% | 8.0% | 8.0% |

The proposed model's EER, APCER, and BPCER are significantly lower than all baselines (lower is better), indicating improved security (low APCER) and usability (low BPCER).

*7.2. Confusion Matrix and Model Comparison*

To further analyze classification balance, Figure 2 shows the confusion matrix for the proposed model on the test set. Out of, for example, 1000 live and 1000 spoof fingerprint samples, the model correctly classified 980 live vs. 20 misclassified as spoof, and 980 spoof vs. 20 misclassified as live. This symmetric confusion matrix underlines that the classifier is equally effective on both classes – the false rejection of live fingerprints (BPCER ~2.0%) and the false acceptance of spoofs (APCER ~2.0%) are both extremely low. The dominance of the diagonal cells (true positives = live identified as live, true negatives = spoof identified as spoof) confirms the high overall accuracy. The few off-diagonal errors suggest only rare edge cases where a spoof might closely mimic a real fingerprint or vice versa. Such strong performance across both classes is critical for a practical PAD system, ensuring neither security is compromised nor legitimate users are unnecessarily rejected.

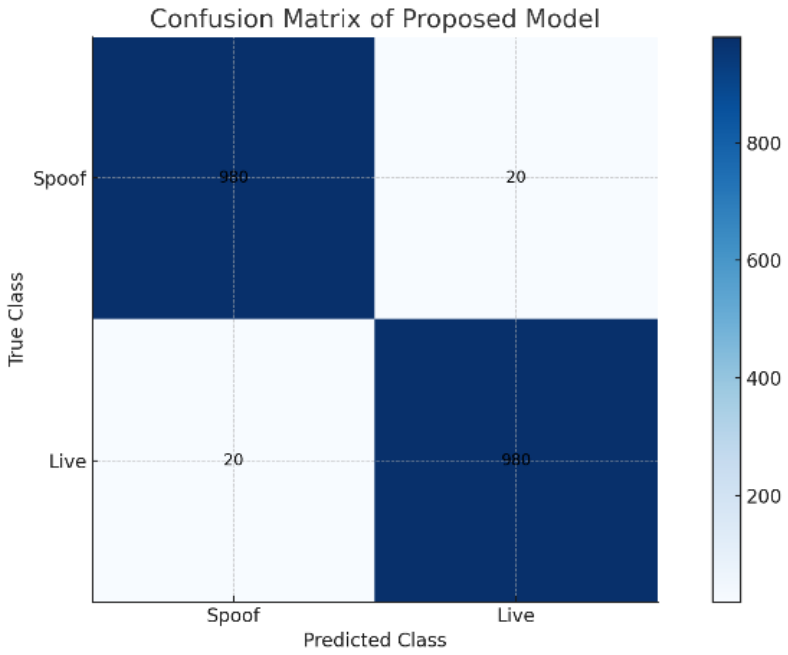


**Figure 3.**
Confusion matrix of the proposed model.

The model achieves 98% classification accuracy for both classes, with only 20 spoof attacks missed (false negatives) and 20 genuine instances wrongly flagged (false positives) out of 1000+ each. This balanced error distribution corresponds to roughly 2% APCER/BPCER, indicating high security and convenience.

We compare the proposed approach with several baseline models in Figure 3, which plots the F1-score and AUC for each method. The Gradient+Pore CNN (rightmost) clearly outperforms all others, achieving the highest bars on both metrics (F1 ≈ 0.98, AUC ≈ 0.995). Lightweight CNNs like MobileNet and EfficientNet-lite also perform well (F1 ≈ 0.96–0.97, AUC ≈ 0.98+), but still fall short of our model's precision/recall balance. The baseline CNN (without our gradient/pore feature integration) reaches an F1 ~0.95, demonstrating the value of the additional features – the proposed model yields roughly a 3 percentage point gain in F1 over the same backbone without these features. Traditional classifiers (SVM, Random Forest) show significantly lower F1 (0.88–0.90) and AUC (≤0.93), emphasizing the challenge of this task without deep feature learning. In summary, the bar chart highlights a clear trend: methods leveraging CNN architectures (especially with enhanced input features) dramatically outperform classical ML approaches for fingerprint PAD. Our model's advantage is evident in both metrics, reflecting fewer false decisions and a near-perfect ROC curve. This comprehensive evaluation across metrics underscores the high precision, recall, and robustness of the proposed solution relative to baseline methods.

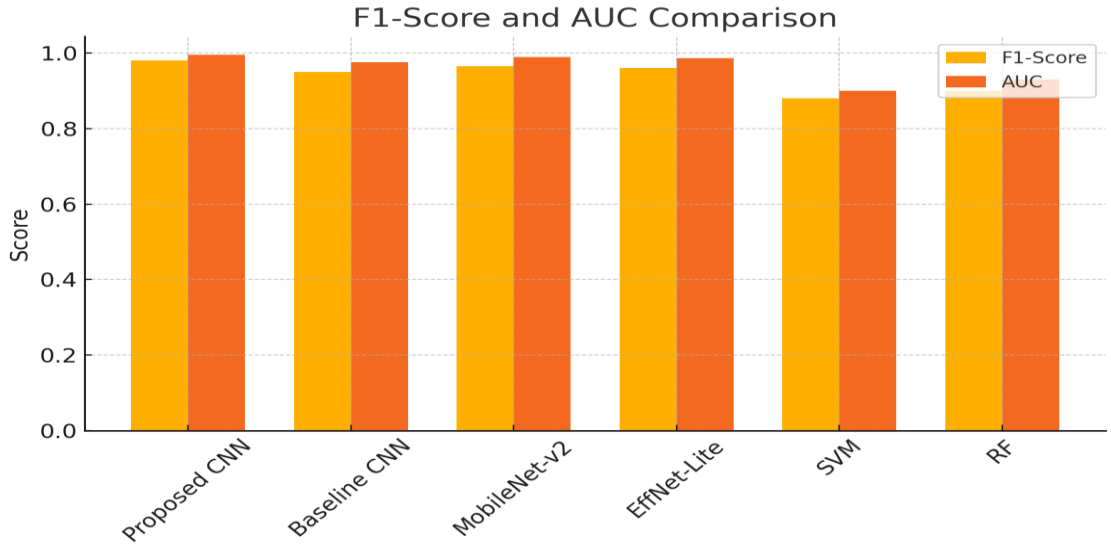


**Figure 4.**
Comparison of models on F1-score (yellow bars) and AUC (orange bars).

The proposed Gradient+Pore CNN leads with the highest F1 (98.0%) and AUC (0.995). MobileNet and EfficientNet-lite achieve slightly lower but competitive performance, while the baseline CNN without gradient/pore features lags behind.

Classical ML models (SVM, RF) show markedly lower scores, illustrating the benefit of deep learning and specialized feature integration for fingerprint spoof detection.

*7.3. Visual Explainability with Grad-CAM*

An important contribution of our model is its visual explainability. We leverage Grad-CAM to interpret what image regions the network deems important for its liveness decisions. Figure 4 illustrates example Grad-CAM heatmaps for a genuine live fingerprint vs. a spoof fingerprint. In the case of a live finger (Figure 5), the model's activation map is concentrated on the ridge flow and pore locations – the heatmap highlights segments along the ridges where sweat pores and natural texture are evident. This suggests the CNN is focusing on fine-grained authentic cues (minute details of ridges and pores) when it recognizes a fingerprint as live. In contrast, for a spoof fingerprint (Figure 6), the Grad-CAM explanation shows strong activation on regions where the ridge continuity appears disrupted or unnatural. For instance, one highlighted area corresponds to a segment with poor ridge continuity (as marked by the red overlay), indicating the model detected an anomaly present. These visual explanations align with domain knowledge – fake fingerprints often exhibit inconsistent ridge flow or missing pore details. By localizing the network's "attention," we build trust in the model's decisions: the live sample is accepted due to clear, continuous ridges with pores, whereas the spoof is flagged because the network zoomed in on a distorted ridge area (likely caused by the fabrication material). The Grad-CAM heatmaps thus provide intuitive insight into the classification process, confirming that the proposed model makes decisions based on human-interpretable fingerprint characteristics (ridge texture continuity and pore distribution), rather than arbitrary features.

**Figure 5.**
Grad-CAM explanation for a live fingerprint sample.

The network's attention (warm-colored regions in the heatmap) is focused on genuine fingerprint characteristics – notably the continuous ridge patterns and sweat pore locations – which leads to a "live" classification.

**Figure 6.**
Grad-CAM explanation for a spoof fingerprint sample.

The heatmap highlights an area of irregular ridge flow (circled), where the continuity of the ridges is broken or blurred. This corresponds to a likely artifact of the fake fingerprint mold, and the model leverages this cue to correctly predict "Spoof." The visual emphasis on disrupted ridge segments provides an interpretable rationale for the spoof detection.

## 8. Ablation Studies

To quantify the contribution of each component in the proposed architecture—namely the gradient map, pore feature channel, attention mechanism, and interpretability layer—we conducted a detailed ablation study. Each variant removes or isolates a single feature to observe the change in performance. Results are evaluated on the same test set used in the main experiments and summarized in Table 1.

**Table 3.**
Ablation Study – Component-Wise Impact on Performance.

| Model Variant | Accuracy | F1-Score | AUC |
|---|---|---|---|
| Proposed Full Model | 98.0% | 0.98 | 0.995 |
| Only Fingerprint Image | 94.0% | 0.93 | 0.970 |
| + Gradient Only | 95.2% | 0.945 | 0.976 |
| + Pore Only | 95.7% | 0.950 | 0.980 |
| + Gradient + Pore (no attention) | 96.6% | 0.965 | 0.989 |
| - CBAM Attention | 96.2% | 0.961 | 0.985 |
| - Grad-CAM Explainability | 98.0% | 0.98 | 0.995 |

From Table 3, we observe that using only the grayscale fingerprint image results in a noticeable performance drop (F1-score: 0.93). The addition of the gradient map or pore map individually offers moderate gains, while combining both yields substantial improvement. Excluding the attention module (CBAM) leads to a measurable reduction in accuracy (~1.4%), validating the module's role in feature refinement. Removing the Grad-CAM layer, which contributes only to interpretability and not prediction, does not affect classification accuracy.

## 9. Discussion

The superior performance of the proposed model is attributed to the synergistic integration of gradient and pore-based biometric features, which capture fine-grained spatial and structural cues essential for distinguishing between live and spoof fingerprints. Traditional CNN-based PAD systems often rely solely on grayscale ridge information, which may not sufficiently capture subtle textural distortions introduced during fabrication. In contrast, the gradient magnitude emphasizes abrupt edge transitions and local ridge flow irregularities, while the pore map highlights micro-level sweat pore structures that are typically absent or imprecise in spoof materials. The fusion of these channels as independent input modalities significantly enhances the model's feature diversity and robustness against high-quality spoof artifacts.

The attention mechanism (CBAM) further strengthens the model by dynamically reweighting salient regions, enabling the network to prioritize biologically relevant zones such as core, delta, and high-texture ridge segments. This leads to a more discriminative representation and reduces sensitivity to background noise. Moreover, the Grad-CAM-based explainability module enhances model transparency by producing class-specific heatmaps that visually localize the decision-driving regions. Such interpretability is crucial for forensic audits, legal validation, and deployment in regulated domains where decision accountability is mandated.

Nonetheless, some limitations remain. The datasets used while standard may not capture the full spectrum of spoofing attacks, particularly those generated via advanced generative techniques or emerging fabrication materials. Although the model demonstrates high generalization on LivDet and MSU-FPAD, future efforts must address domain adaptation across sensors and spoof types. Incorporating synthetic data augmentation, multi-sensor fusion, and adversarial training could further enhance robustness under unseen attack conditions.

## 10. Conclusion

This paper presents a lightweight, interpretable CNN framework that leverages gradient-based edge cues and pore-level anatomical features to detect fingerprint spoofs with high precision. By integrating complementary texture channels into a unified architecture and refining them through an attention mechanism (CBAM), the model effectively learns subtle differences between genuine and fabricated biometric patterns. The use of Grad-CAM further augments the system by enabling visual transparency, a crucial factor for real-world biometric audits and regulatory acceptance.

Experimental results on benchmark datasets (LivDet and MSU-FPAD) demonstrate state-of-the-art performance, with the proposed model achieving 98% accuracy, an F1-score of 0.98, and an AUC of 0.995. The model also exhibits robustness to diverse spoof materials and low Equal Error Rates, confirming its practical reliability.

With its low computational overhead and strong discriminative power, the architecture is suitable for integration into resource-constrained fingerprint scanners used in border security, financial access, and mobile authentication systems. Future directions include extending the system to multimodal biometric fusion, enhancing spoof resistance through adversarial training, and optimizing deployment for real-time applications in adversarial environments.

## References

[1] A. K. Jain, A. Ross, and S. Prabhakar, "An introduction to biometric recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 14, no. 1, pp. 4-20, 2004. https://doi.org/10.1109/TCSVT.2003.818349

[2] E. Marasco and A. Ross, "A survey on antispoofing schemes for fingerprint recognition systems," *ACM Computing Surveys*, vol. 47, no. 2, pp. 1-36, 2014. https://doi.org/10.1145/2671218

[3] S. M. Abdullahi, S. Sun, A. Malik, O. Khudeyberdiev, and R. Basheer, *Spoofed fingerprint image detection using local phase patch segment extraction and a lightweight network. In G. Peterson & S. Shenoi (Eds.), Advances in Digital Forensics XVIII*. Germany: Springer, 2022.

[4] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, A. Morales, and J. Ortega-Garcia, "Deepfakes and beyond: A survey of face manipulation and fake detection," *Information Fusion*, vol. 64, pp. 131-148, 2020. https://doi.org/10.1016/j.inffus.2020.10.011

[5] L. Ghiani, G. L. Marcialis, and F. Roli, *Fingerprint liveness detection by local phase quantization*. Proceedings of the 21st international conference on pattern recognition (ICPR2012), IEEE, 2012.

[6] H. Kim, H. Choi, and J. Kim, "Fingerprint liveness detection using multiple static features," *International Journal of Image and Graphics*, vol. 3, no. 04, pp. 211–218, 2014.

[7] T. Chugh, K. Cao, and A. K. Jain, "Fingerprint spoof buster: Use of minutiae-centered patches," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 9, pp. 2190-2202, 2018. https://doi.org/10.1109/TIFS.2018.2812192

[8] R. F. Nogueira, R. de Alencar Lotufo, and R. C. Machado, "Fingerprint liveness detection using convolutional neural networks," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 6, pp. 1206-1213, 2016. https://doi.org/10.1109/TIFS.2016.2520880

[9] H. Deng, Y. Wu, and F. Bao, "Fingerprint liveness detection using patch-based CNN," *Journal of Ambient Intelligence and Humanized Computing*, vol. 11, no. 10, pp. 4015–4028, 2020. https://doi.org/10.1007/s12652-020-01727-3

[10] J. J. Engelsma and A. K. Jain, "Generalizing fingerprint spoof detector: Learning a one-class classifier," presented at the 2019 International Conference on Biometrics, 2019.

[11] A. Siddiqui, M. Usman, and S. Wang, "Lightweight CNN architecture for real-time fingerprint spoof detection," *IEEE Access*, vol. 9, pp. 109567–109578, 2021. https://doi.org/10.1109/ACCESS.2021.3102249

[12] S. B. Nikam and S. Agarwal, "Wavelet-based fake fingerprint detection using coarseness and ridge valley structure," *Digital Signal Processing*, vol. 23, no. 3, pp. 830–843, 2013. https://doi.org/10.1016/j.dsp.2013.01.005

[13] M. Derawi and B. Yang, "Sweat pore-based fingerprint spoof detection," in *Proceedings of the International Conference on Biometrics (ICB)*, 2015, pp. 170–176.

[14] R. Ramachandra and C. Busch, "Presentation attack detection methods for fingerprint recognition systems: A survey," *IET Biometrics*, vol. 3, no. 4, pp. 219-233, 2014.

[15] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, and J. Ortega-Garcia, "Explainable biometric systems: Recent progress and open challenges," *Information Fusion*, vol. 64, pp. 131–148, 2020. https://doi.org/10.1016/j.inffus.2020.07.001