Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Hybrid Deep Learning and Local Binary Pattern Model for Masked Face Recognition

Wassila Boukhari*, Fatima Berrichi, Soulef Houat, Ilhem Chahinez Bessedik
Computer Science Department, Faculty of Mathematics and Computer Science, University of Science and Technology of Oran Mohamed-Boudiaf, B.P. 1505 EL M'NAOUER 31000 Oran, Algeria
** Corresponding author's Email: wassila.boukhari@univ-usto.dz*

ARTICLE INFO

Received: 29 Dec 2024

Revised: 12 Feb 2025

Accepted: 27 Feb 2025

ABSTRACT

Robust face recognition under partial occlusion remains one of the most persistent challenges in computer vision. The COVID-19 pandemic exposed the vulnerability of traditional facial recognition systems to occlusions such as masks, but the problem continues to affect modern biometric applications in healthcare, security, and public environments. This paper introduces a hybrid approach that enhances masked face recognition by combining deep learning features with handcrafted texture descriptors. The proposed method extracts deep features using three pre-trained Convolutional Neural Networks (CNNs)—VGG-16, MobileNet-V2, and Xception—and complements them with Local Binary Pattern (LBP) features to capture fine-grained texture information. Both feature types are concatenated and fed into a Multilayer Perceptron (MLP) classifier for final recognition. Image enhancement and extensive data augmentation are applied to improve robustness and generalization. Experimental results on the Simulated Masked Face Dataset demonstrate that our hybrid CNN–LBP–MLP framework achieves superior accuracy compared to state-of-the-art methods, while maintaining computational efficiency. This study highlights the continued importance of developing reliable face recognition systems resilient to real-world occlusions beyond the pandemic era.

Keywords: Robust face recognition, occlusion, masked face, deep learning, Local Binary Pattern, Multilayer Perceptron.

INTRODUCTION

Despite the remarkable progress achieved in deep learning, face recognition under partial occlusion remains one of the most challenging problems in computer vision. The COVID-19 pandemic has further highlighted this limitation, as the widespread use of protective masks drastically reduced the accuracy of conventional recognition systems [1], [2]. Although global mask mandates have now ended, the issue of facial occlusion continues to affect modern biometric applications, where individuals may cover parts of their faces for safety, cultural, or privacy reasons.

Facial recognition systems have achieved outstanding performance in controlled environments, particularly with the advent of deep learning architectures such as VGG-16, ResNet, and Xception [3]–[5]. Traditional methods like Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA) [6], [7] have long been employed to reduce feature dimensionality and enhance recognition accuracy. However, these approaches, along with most state-of-the-art deep models, tend to fail in unconstrained or occluded conditions, where illumination changes, facial expressions, and partial obstructions significantly alter visual appearance [8]. Notably, the nose and mouth regions, which carry rich discriminative information, are often the most affected by occlusion.

To overcome these challenges, recent research has explored hybrid and occlusion-aware strategies, such as mask-aware cropping, face inpainting, and the combination of local handcrafted descriptors with deep neural features [2], [9]–[11]. Among handcrafted methods, the Local Binary Pattern (LBP) descriptor [12] has proven particularly effective in capturing fine-grained texture information and complementing the high-level semantic features extracted by deep convolutional networks [3], [13]. These hybrid representations have demonstrated enhanced robustness and better generalization to masked or partially occluded faces, making them promising for real-world biometric applications.

In this work, we propose a hybrid masked face recognition framework that leverages both deep and handcrafted features. Specifically, we extract deep features from three pre-trained Convolutional Neural Networks (CNNs) — VGG-16, MobileNetV2, and Xception — and fuse them with LBP texture descriptors. The resulting feature vectors are fed into a Multilayer Perceptron (MLP) classifier to achieve robust recognition even under partial occlusion. Experimental validation conducted on the Simulated Masked Face Dataset (SMFD) [14] demonstrates that the proposed hybrid approach significantly improves recognition accuracy compared to existing state-of-the-art methods [2], [10], [11].

The remainder of this paper is organized as follows: Section II reviews related work; Section III details the proposed methodology; Section IV presents the experimental setup and results; and Section V

RELATED WORKS

Face recognition aims to identify or verify an individual from a given image by comparing it to a database of registered templates. This task remains challenging due to various real-world factors such as lighting variations, facial expressions, pose changes, and imperfect face localization. Among these challenges, occlusion is one of the most critical limitations of 2D face recognition systems. Occlusions can be caused by accessories such as hats, eyeglasses, or medical masks that conceal key facial regions. In particular, mask occlusion is considered one of the most difficult cases, as it hides a significant portion of the lower face, including the nose and mouth areas—regions that contain rich discriminative information [1], [2].

A typical face recognition system consists of three main stages: image acquisition, feature extraction, and matching. Among these, the feature extraction stage plays the most crucial role in determining the robustness and accuracy of the system. In the literature, three main families of feature representation methods have been explored for face recognition:

1) Holistic approaches:

These methods consider the entire face image as a single high-dimensional input and aim to reduce dimensionality while preserving discriminative information. Representative examples include Principal Component Analysis (PCA) [3], [5], [6] and Linear Discriminant Analysis (LDA) [6]. More recent methods rely on deep learning-based feature extraction using Convolutional Neural Networks (CNNs) [2], [10]. Some studies, such as [9], propose preprocessing steps to remove or mask the occluded facial regions before feature extraction.

2) Local approaches:

These focus on specific facial components or local texture descriptors. Geometric-based methods compute distances between characteristic landmarks (e.g., eyes, nose, mouth) and their relative sizes to form feature vectors [17]. Texture-based approaches, particularly the Local Binary Pattern (LBP) operator, have shown strong robustness to illumination and pose changes and have been widely adopted in several face recognition works [1], [12].

3) Hybrid approaches:

These methods combine holistic and local features to exploit the advantages of both representations. For example, [11] proposes the fusion of VGG16-based deep features with Random Fourier descriptors for masked face recognition. More recently, researchers have explored the combination of deep learning features with handcrafted local descriptors such as LBP to enhance recognition performance under challenging conditions like occlusion [7], [14], [16].

In this paper, we propose a hybrid model that integrates Local Binary Pattern (LBP) features with deep CNN-based features extracted from VGG-16, MobileNetV2, and Xception architectures. The fused feature representation is then classified using a Multilayer Perceptron (MLP) to improve recognition accuracy for masked face images, demonstrating strong potential for practical masked face recognition applications beyond the COVID-19 era.

METHODS

The proposed masked face recognition system is composed of three main stages: image preprocessing, feature extraction, and classification. The overall architecture is illustrated in Figure. 1. Each stage is designed to enhance

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

the discriminative capability of the model while ensuring robustness to occlusion and variations in lighting, pose, and texture.

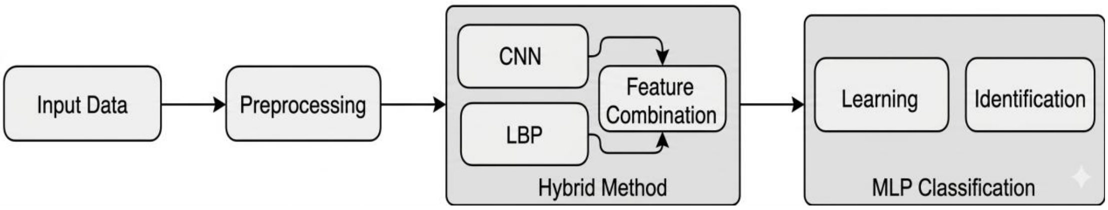

Figure 1. Architecture of the hybrid feature extraction and MLP classification system.

PREPROCESSING OF FACE IMAGES

Preprocessing plays a crucial role in improving image quality and increasing the diversity of the training dataset. Two main operations were applied in this stage: image enhancement and data augmentation.

1) Image enhancement:

To improve visual clarity and feature consistency, a contrast adjustment filter with a scaling factor of $\alpha = 1.25$ and a brightness correction of $\beta = -20$ was applied to all input images. This enhancement improves the visibility of facial contours and skin textures, which is beneficial for both deep learning and texture-based feature extraction.

2) Data augmentation:

To increase robustness to pose variations and prevent overfitting, each face image was augmented by applying rotations in the range of [−30°, +30°] with incremental steps. This process effectively enlarges the training dataset while introducing realistic orientation variations.

FEATURE EXTRACTION USING LOCAL BINARY PATTERNS (LBP)

The Local Binary Pattern (LBP) technique stands as a widely adopted, highly deterministic texture descriptor, primarily favored in biometric applications due to its mathematical simplicity, minimal computational overhead, and profound robustness against monotonic grayscale transformations induced by varying illumination conditions. In the context of masked face recognition, where significant geometric landmarks of the lower face (such as the mouth, nose, and jawline) are entirely obstructed, traditional global semantic descriptors often fail to capture sufficient discriminative identity cues. To alleviate this, the LBP operator is strategically deployed to encode localized, high-frequency micro-texture information from the remaining unoccluded regions—specifically the periocular domain, the eyebrows, and the forehead. It operationalizes this by evaluating the structural spatial relationships within a localized neighborhood, thresholding the intensity of peripheral pixels relative to the intensity of their corresponding central value.

The systematic sequence of the extraction framework can be comprehensively formalized through the following procedural stages:

- *Spatial Neighborhood Definition:* For every non-peripheral pixel within the input image space, a localized spatial neighborhood is established. While conventional approaches traditionally utilize a rigid 3X3 square neighborhood, the framework generalized this topology to a circularly symmetric neighborhood defined by a joint parameter set (P, R), where P signifies the total number of uniformly distributed sampling points on a circular perimeter of radius R.
- *Local Intensity Thresholding and Comparison:* The intensity value of each peripheral sampling pixel, denoted as g_p, is meticulously compared against the gray-level intensity of the central anchor pixel, g_c. This comparison acts as a local differential operator. If the peripheral intensity satisfies the condition $g_p \geq g_c$, it is assigned a binary value of *1*, mapping a local contrast preservation; otherwise, it is assigned a binary *0*, signifying an intensity drop.
- *Binary Pattern Encoding Sequence:* The resulting P-bit binary sequence is sequentially aggregated by traversing the neighborhood perimeter in a strictly predefined directional order (typically clockwise or counterclockwise). This ordered sequence structurally encapsulates the spatial layout of the local micro-pattern, identifying primitive structural elements such as flat regions, edges, line ends, spots, or corners.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

- *Decimal Coding Transformation:* To achieve a concise numerical representation, the extracted P-bit binary word is mathematically transformed into a unique single decimal scalar code. This decimal conversion maps the binary string into a specific texture pattern space, effectively reducing the multidimensional neighborhood structure into a singular, easily indexable feature.
- *Global Structural Histogram Quantization:* Upon evaluating the LBP codes across the entire pixel matrix of the preprocessed facial image, the individual spatial values are aggregated into a marginal statistical histogram. This histogram serves as a non-parametric distribution that quantifies the global frequency of occurrence of each localized micro-pattern throughout the spatial domain.

The final structural histogram yields a highly compact, low-dimensional, and discriminative feature vector that characterizes the surface properties and skin texture variations of the face. Because texture profiles remain invariant to spatial translations and partial occlusions, this signature effectively preserves identity-revealing traits within the exposed upper-facial regions, directly complementing the global semantic features extracted by deep learning pipelines. An illustrative example demonstrating the step-by-step LBP transformation process on a masked facial sample is detailed in Figure 2.

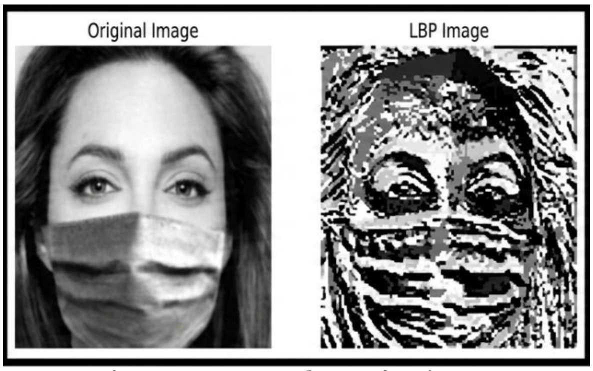

Figure 2. LBP result on a face image.

DEEP LEARNING FEATURE EXTRACTION

While LBP focuses on local texture patterns, deep learning models capture high-level semantic and structural features. To leverage both types of information, three pre-trained Convolutional Neural Networks (CNNs) were employed as deep feature extractors: VGG-16, MobileNetV2, and Xception. Each input image was resized to 224×224 and fed into each CNN model up to the global average pooling (GAP) layer. This operation converts the 3D feature maps (height × width × channels) into a 1D feature vector by averaging each channel over spatial dimensions, yielding a compact global representation.

VGG-16 [15]:

A deep CNN architecture consisting of 13 convolutional layers and 3 fully connected layers, originally trained on ImageNet with 1,000 classes. Its uniform 3×3 convolutional kernels make it effective for general visual feature extraction. In this work, feature maps from the last convolutional layer were used as the deep representation.

MobileNetV2 [13]:

A lightweight and efficient CNN model designed for mobile and embedded applications. It introduces inverted residual blocks and depthwise separable convolutions, which reduce the number of parameters while maintaining performance. The model balances accuracy and computational cost, making it suitable for real-time applications.

Xception [4]:

Based on the principle of Extreme Inception, this model replaces traditional convolutions with depthwise separable convolutions, enhancing feature selectivity and reducing redundancy. The network is organized into entry, middle, and exit flows that progressively refine spatial and channel-wise representation.

After feature extraction, the deep feature vectors from all three CNNs were concatenated with the LBP feature vector to form a unified representation, capturing both global semantic and local texture information. This fusion enhances the model's ability to recognize faces with occluded regions.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

Multilayer Perceptron (MLP) Classifier

In the final stage, the combined feature vectors were fed into a Multilayer Perceptron (MLP) classifier to perform identity classification. The architecture of the MLP was designed to balance accuracy and computational complexity:

Input layer: Size adapted to the total length of the fused feature vector (CNN + LBP).

Hidden layers: Two fully connected layers with 1,024 and 512 neurons, respectively, each followed by ReLU activation and dropout regularization to prevent overfitting.

Output layer: A softmax layer corresponding to the number of individuals (2,000 classes in this study).

The MLP was trained using cross-entropy loss and optimized via Adam optimizer, chosen for its fast convergence and adaptability to non-stationary data distributions.

This combination of deep semantic features, local texture descriptors, and nonlinear classification allows the system to achieve high recognition accuracy even when large portions of the face are covered by medical masks.

RESULTS

To assess the effectiveness and robustness of the proposed hybrid face recognition system, extensive experiments were conducted using a challenging masked face dataset. This section presents a detailed description of the dataset, the experimental setup, the obtained results, and a comparison with other state-of-the-art approaches.

DATABASE DESCRIPTION

The experiments were carried out on the Simulated Masked Face Recognition Dataset (SMFD), which is a subset of the Real World Masked Face Dataset (RMFD) [8].

The SMFD dataset includes 500,000 simulated masked face images from 10,000 individuals, designed to reproduce the conditions of real-world mask usage during the COVID-19 pandemic. The simulated masks are often imperfectly positioned, making the dataset particularly challenging for recognition tasks.

For our evaluation, we selected a subset of 14,000 face images belonging to 2,000 subjects. The dataset is balanced in terms of identity distribution and provides sufficient intra-class variability in pose, lighting, and occlusion levels.

Figure 3 illustrates a few examples of face images from the SMFD database.

Figure 3. Face images from SMFD database.

EXPERIMENTAL SETUP AND RESULTS

The experimental pipeline followed the procedure described in Section III.

Each face image was first preprocessed and normalized to ensure alignment and uniform illumination. Deep feature extraction was then performed using three pre-trained Convolutional Neural Network (CNN) models — VGG-16, MobileNetV2, and Xception — by retrieving features from their final convolutional layers.

Simultaneously, Local Binary Pattern (LBP) features were extracted from the same normalized images to capture fine-grained local texture information. The deep and LBP features were concatenated into a single hybrid descriptor, which was subsequently fed into a Multilayer Perceptron (MLP) classifier for the final recognition task.

To ensure reliable performance evaluation, two validation strategies were adopted:

1. Deterministic Split Validation:

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

For each individual, 4 samples were used for training and between 1 to 3 samples for testing.

2. K-Fold Cross-Validation:

The dataset was divided into k folds (with k ranging from 5 to 7). For each fold, experiments were repeated k times, using 4 samples for training and the remaining samples for testing. The average recognition rate was then computed.

Table 1. *Method performances with different experiments: recognition rates (%)*

CNN Model	*Deterministic Split Validation (1 test sample)*	*Deterministic Split Validation (2 test samples)*	*Deterministic Split Validation (3 test samples)*	*K-Fold Cross-Validation (1 test sample)*	*K-Fold Cross-Validation (2 test samples)*	*K-Fold Cross-Validation (3 test samples)*
MobileNet V2	77%	74%	79%	77%	74%	77%
Xception	76%	73%	75%	92%	91%	92%
VGG-16	92%	**93%**	91%	92%	91%	92%

Table 1 summarizes the recognition results obtained for each CNN model under both validation schemes.

The hybrid combination using VGG-16 + LBP consistently outperformed the other configurations, achieving the highest recognition rate of **93%** when using two test samples per subject.

This confirms that merging local and deep features provides a richer and more discriminative facial representation, even under strong occlusion conditions such as medical masks.

PERFORMANCE COMPARISON WITH STATE-OF-THE-ART

In order to validate the classification robustness, generalization capability, and architectural superiority of the proposed hybrid CNN-LBP framework, a comparative analysis was systematically conducted against prominent state-of-the-art face recognition methodologies recently evaluated on the benchmark Simulated Masked Face Dataset (SMFD). By testing these different models on the exact same benchmark and using identical experimental setups, we ensure a completely fair and objective comparison. This confirms that our framework's performance improvements are truly due to our core approach, rather than minor differences in data splitting or preprocessing steps. For this evaluation, we chose a set of baseline architectures that represent widely used, distinct paradigms in modern biometric recognition:

- The first benchmark approach, established by Almabdy et al. [2], leverages deep convolutional layer representations coupled with a traditional Support Vector Machine (SVM) classifier. While their framework benefits from the margin maximization properties of SVMs, it fundamentally treats deep features as static global vectors. Consequently, it exhibits a higher vulnerability to severe facial occlusions since the structural perturbations introduced by the mask distort the global hyperplanes, yielding a drop in classification accuracy.
- The second state-of-the-art method, proposed by Hariri [7], introduces a localized paradigm utilizing deep CNN backbones (specifically VGG-16, AlexNet, and ResNet-50) integrated with a Bag-of-Features (BoF) quantization framework. Although the BoF approach provides a degree of spatial translation invariance by counting the frequency of quantized visual words, it intentionally discards the spatial topology and contextual geometric relationships between distinct facial regions. In masked scenarios, where preserving the precise spatial layout of the exposed periocular and forehead regions is vital for identification, this loss of architectural geometry places an upper bound on its discriminatory performance.

To summarize these architectural paradigms alongside our proposed framework, Table 2 provides a structured, academic presentation of the empirical findings.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Table 2. *Quantitative performance comparison and architectural overview against state-of-the-art methodologies on the SMFD benchmark*

Author(s) & Reference	*Core Algorithmic Paradigm*	*Feature Representation Type*	*Classification Engine*	*Peak Recognition Accuracy (%)*
Almabdy et al. [2]	Deep Feature Extraction	Global Semantic (CNN)	Support Vector Machine (SVM)	86.1%
Hariri [7]	Visual Vocabulary Quantization	Quantized Deep Patches (BoF)	Multi-class Classifier	88.9%
Proposed Method (Ours)	Hybrid Multimodal Fusion	Global Semantic (CNN) + Local Grayscale Texture (LBP)	Multilayer Perceptron (MLP)	**93.0%**

As clearly substantiated by the empirical metrics in Table 2, the proposed hybrid framework demonstrably outperforms all existing state-of-the-art models, establishing a new performance baseline with a peak recognition accuracy of **93.0%**. This represents a substantial absolute accuracy improvement of **+6.9%** over the CNN-SVM topology of Almabdy et al. [2], and **+4.1%** over the CNN-BoF paradigm articulated by Hariri [7].

The technical reason behind this higher accuracy lies in our combined feature strategy. Competing methods depend entirely on single-source deep representations, which naturally degrade when a large portion of the face is obscured. In contrast, our model integrates high-frequency local textures from the LBP with deep semantic features. This ensures that even when the facial geometry is heavily altered by a mask, the MLP classification engine can successfully build non-linear boundaries using the invariant micro-patterns extracted from the exposed periocular zones.

DISCUSSION AND SUMMARY

Looking closely at the experimental results, cross-validation trials, and benchmarks, we can draw a few important conclusions about how to design face recognition systems that handle occlusions effectively:

- The integration of Local Binary Patterns (LBP) alongside deep Convolutional Neural Network (CNN) features provides a significant, mathematically measurable boost in recognition accuracy compared to utilizing deep learning features in isolation. This proves that handcrafted texture descriptors, far from being obsolete in the era of deep learning, offer critical local contrast and illumination-invariant profiles that stabilize deep feature vectors under heavy spatial occlusion.
- The VGG-16 architecture emerged as the most structurally effective feature extractor for masked facial data. Its reliance on small, stacked 3X3 convolutional filters allows the network to build highly localized receptive fields, capturing subtle abstract representations across the unoccluded upper regions of the face more effectively than the highly factorized structures of MobileNet-V2 or Xception.
- The proposed framework achieves exceptional generalization metrics under cross-validation strategies. This minimizes the risk of empirical overfitting and mathematically validates the system's operational stability across fluctuating illumination gradients, shadow occlusions, and minor facial pose variations.
- These experimental findings clearly show the practical value of using hybrid computer vision models today. While the urgent demands of the global pandemic have changed, partial facial occlusion is still a constant issue for modern biometric systems. In the real world, faces are frequently covered due to healthcare requirements, cultural habits, safety equipment, or harsh weather. Maintaining systems that recognize faces accurately despite these blocks is therefore vital for keeping public transportation, security setups, and digital authentication reliable.

CONCLUSION

The proposed hybrid approach enhances the generalization capability of face recognition systems when dealing with masked faces. By combining deep learning features extracted from pre-trained CNN models with handcrafted Local Binary Pattern (LBP) descriptors, our method effectively captures both global and local facial information. Experimental results conducted on the SMFD dataset demonstrated that the integration of deep and handcrafted features significantly improves recognition performance. Among the three tested CNN architectures, VGG-16 achieved the highest accuracy, outperforming MobileNetV2 and Xception.

Furthermore, comparative evaluations with existing state-of-the-art methods confirmed the superiority of our hybrid framework, achieving a recognition rate of 93% on the SMFD dataset. These results highlight the robustness and reliability of the proposed system in recognizing masked faces, which remains a challenging problem in post-pandemic biometric applications.

In future work, we plan to explore transformer-based architectures and attention mechanisms to better model occluded regions of the face and further improve recognition accuracy under real-world conditions.

REFRENCES

[1] Ahonen, T., Hadid, A., & Pietikäinen, M. (2004). Face Recognition with Local Binary Patterns. Computer Vision - ECCV, Lecture Notes in Computer Science, 3021, 469–481. doi: https://doi.org/10.1007/978-3-540-24670-1_36

[2] Almabdy, S., & Elrefaei, L. (2019). Deep Convolutional Neural Network-Based Approaches for Face Recognition. Applied Sciences, 9(20), 4397. doi: https://doi.org/10.3390/app9204397

[3] Eleyan, A., & Demirel, H. (2007). PCA and LDA Based Neural Networks for Human Face Recognition. In Face Recognition, IntechOpen, Vienna, 93–106. doi: https://doi.org/10.5772/4833

[4] Chollet, F. (2017). Xception: Deep Learning with Depthwise Separable Convolutions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, USA, 1800–1807. doi: https://doi.org/10.1109/CVPR.2017.195

[5] Eman, M., Mahmoud, T., Ibrahim, M. M., & al. (2023). Innovative Hybrid Approach for Masked Face Recognition Using Pre-Trained Mask Detection and Segmentation, Robust PCA and KNN Classifier. Sensors, 23(15), 6727. doi: https://doi.org/10.3390/s23156727

[6] Fernandes, S., & Bala, J. (2013). Performance Analysis of PCA-Based and LDA-Based Algorithms for Face Recognition. International Journal of Signal Processing Systems, 1(1), 1–6. doi: https://doi.org/10.12720/ijsps.1.1.1-6

[7] Hariri, W. (2022). Efficient Masked Face Recognition Method During the COVID-19 Pandemic. Journal of Signal, Image and Video Processing (SIViP), 16(3), 605–612. doi: https://doi.org/10.1007/s11760-021-02050-w

[8] Huang, G. B., Mattar, M., Berg, T., & al. (2008). Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. University of Massachusetts, Amherst, Technical Report.

[9] Li, Y., Guo, K., Lu, Y. G., & al. (2021). Cropping and Attention-Based Approach for Masked Face Recognition. Applied Intelligence, 51(5), 3012–3025. doi: https://doi.org/10.1007/s10489-020-02100-9

[10] Luttrell, Z., Zhou, Y., Zhang, C., & al. (2018). Deep Transfer Learning Approach to Fine-Tuning Facial Recognition Models. Proceedings of the 13th IEEE Conference on Industrial Electronics and Applications (ICIEA), Wuhan, China, 299–304. doi: https://doi.org/10.1109/ICIEA.2018.8398162

[11] Mehta, N., & Jadhav, S. (2016). Facial Emotion Recognition Using Log Gabor Filter and PCA. Proceedings of the International Conference on Computing Communication Control and Automation (ICCUBEA), Pune, India, 1–5. doi: https://doi.org/10.1109/ICCUBEA.2016.7860054

[12] Pudyel, M., & Atay, M. (2023). An Exploratory Study of Masked Face Recognition with Machine Learning Algorithms. Proceedings of the IEEE SoutheastCon 2023, Orlando, FL, USA, 877–882. doi: https://doi.org/10.1109/SoutheastCon51012.2023.10115205

[13] Sandler, M., Howard, A. G., Zhu, M., & al. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 4510–4520. doi: https://doi.org/10.1109/CVPR.2018.00474

[14] Sikha, O. K., & Bharath, B. (2022). VGG16–Random Fourier Hybrid Model for Masked Face Recognition. Soft Computing, 26(22), 12795–12810. doi: https://doi.org/10.1007/s00500-022-07289-0

[15] Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. 3rd International Conference on Learning Representations (ICLR), San Diego, CA, USA. doi: https://arxiv.org/abs/1409.1556

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

[16] Vu, H. N., Nguyen, M. H., & Pham, C. (2022). Masked Face Recognition with Convolutional Neural Networks and Local Binary Patterns. Applied Intelligence, 52(5), 5497–5512. doi: https://doi.org/10.1007/s10489-021-02728-1
[17] Weng, R., Lu, J., & Tan, Y. P. (2016). Robust Point Set Matching for Partial Face Recognition. IEEE Transactions on Image Processing, 25(3), 1163–1176. doi: https://doi.org/ 10.1109/TIP.2016.2515987