Volume 5, Issue X, October, 2026; No. 92, pp. 1196-1207
Submitted 12/09/2026; Final peer review 01/11/2026
Online Publication 04/11/2026
Available Online at http://www.ijortacs.com

Facial Classification for Pensioner Verification Using Gradient Features and a Feed Forward Neural Network

Aniekwe Stephen Elochukwu
Department of Computer Science, Alex Ekwueme Federal University
Email: aniekwestephen3@gmail.com

Abstract

Reliable pensioner verification requires a clear connection between facial evidence and an enrolled identity. This study examines a gradient-feature and feed-forward neural-network pipeline developed for pension administration in Enugu North, Nigeria. Facial images were acquired, processed by Gaussian smoothing, represented using histograms of oriented gradients, and analyzed within a MATLAB-based workflow. The collection description records 4,030 images from 2018–2021; the available development confusion matrices contain 699 classified observations, partitioned into 489 training, 105 validation and 105 test cases. Reanalysis of these matrices yielded accuracies of 96.93%, 96.19% and 97.14%, respectively. On the test partition, Class 2 precision, recall and F1-score were 94.87%, 97.37% and 96.10%; balanced accuracy was 97.19%. Best validation cross-entropy was 0.041598 at epoch 8, while the recorded run stopped at epoch 14. A separate application table contained 18 True-coded outcomes in 20 cases. The results indicate effective discrimination within the documented two-class experiment, with similar performance across development partitions. The study formalizes the processing model and establishes a count-based evaluation that separates model discrimination from operational identity verification. Wider deployment requires clarification of class semantics, linkage between the collection and development subsets, and independently adjudicated identity claims.

Keywords: **facial classification, pensioner verification, histogram of oriented gradients, neural network, biometric evaluation**

1 Introduction

Civil-service pension administration depends on connecting an individual presenting for verification to a valid personnel record. A name or document identifier supports this connection, but the organization must also establish that the claimant is the person associated with that record. Face-based processing offers a contactless source of identity evidence that can be incorporated into an administrative workflow. The engineering problem is to extract useful information from a captured image and interpret the resulting decision at the level actually evaluated.

Biometric identification searches an enrolled population, whereas verification evaluates a claimed identity against its reference (Jain et al., 2004). A two-class decision can support a verification system, but it is not equivalent to identifying every person within a large gallery. This distinction is particularly relevant when a development experiment reports only two plotted output classes while the acquisition description concerns thousands of images. The number of collected images, the number of model observations and the number of identities therefore require separate treatment.

The Enugu North study combined conventional image processing with neural classification and a MATLAB interface. Its original evaluation figures contain more information than the headline accuracy alone: separate training, validation and test matrices, classwise errors, and a learning curve with a marked validation minimum. These records permit an evaluation of class discrimination, learning behavior and sample uncertainty without inventing additional observations. They also provide a basis for explaining where the recognition component fits within pension administration.

This paper addresses two questions. First, how can the acquisition, gradient representation and neural computation

Corresponding Author Tel: +234 80 6707 3213

be expressed as a coherent mathematical pipeline? Second, what do the available class counts establish about performance across the recorded partitions? The contributions are (1) a mathematically specified Gaussian–gradient–neural classification model for the pensioner-verification application, and (2) a reconstructed partition-level evaluation including precision, recall, F1-score and balanced accuracy. The objective is application-specific model exposition and evidence-based evaluation rather than a claim of a new neural architecture.

2 Literature Review

2.1 Representation and classification

Face recognition methods differ in how they represent an image and how they learn separation between classes. Wang and Deng (2021) review the progression toward deep facial representations and the associated importance of datasets and evaluation protocols. In a feature-based pipeline, the image descriptor is specified before classifier training. This makes the transformation easier to describe and can suit a prototype with modest computing resources, although it does not establish equivalence to modern embedding systems.

Gaussian convolution is an established smoothing operation whose behavior depends on scale, support and boundary treatment (Getreuer, 2013). It can suppress small intensity fluctuations, but increasing the scale also removes detail. Histograms of oriented gradients summarize local gradient distributions and normalize neighboring histograms to reduce sensitivity to contrast changes (Dalal & Triggs, 2005). The descriptor was introduced for human detection; its use here is an application to facial representation, which must be assessed using the present data rather than the original detection benchmark.

2.2 Modern comparison points and research gap

FaceNet learns a compact embedding, while ArcFace introduces an angular-margin classification objective (Schroff et al., 2015; Deng et al., 2019). MagFace and AdaFace subsequently incorporate image quality into representation learning or the training margin (Meng et al., 2021; Kim et al., 2022). These approaches identify relevant future comparators, but none was trained and tested on the present partitions. Their published results are therefore not used as numerical baselines.

Recognition research also distinguishes image matching from presentation-attack detection and demographic evaluation (Grother et al., 2019; Sharma & Selwal, 2023). A classifier may separate the supplied classes while remaining vulnerable to poor captures, unknown identities or replayed photographs. The practical gap addressed here is the absence of a unified mathematical and numerical account of the local prototype. Table 1 summarizes the methodological relationship between the literature and the present investigation.

Table 1: *Relationship of prior methods to the present model*

Research strand	Representative source	Role in this study
Biometric task definition	Jain et al. (2004)	Separate identification, verification and class discrimination
Image filtering	Getreuer (2013)	Explain the Gaussian processing stage
Gradient representation	Dalal and Triggs (2005)	Define the feature vector supplied to the classifier
Embedding learning	Schroff et al. (2015); Deng et al. (2019)	Motivate future same-data comparisons
Quality-aware learning	Meng et al. (2021); Kim et al. (2022)	Motivate quality control and robustness testing

3 Methodology

3.1 Study setting and evaluation records

The study concerned pensioner verification in Enugu North, Nigeria. The acquisition account describes 4,030 facial images collected between 1 January 2018 and 31 December 2021, associated with bearer names and acquired at 240 × 120 pixels. The implementation used MATLAB and its computer-vision and neural-network facilities. Image count is retained separately from person count because the number of unique identities and repeated captures is not specified.

The plotted development experiment contains 699 observations: 489 training, 105 validation and 105 test

Corresponding Author Tel: +234 80 6707 3213

observations. These counts correspond to approximately 69.96%, 15.02% and 15.02% of the plotted experiment. They were obtained by summing the four cells in each original confusion matrix, not by assuming how all 4,030 collected images were divided. The mapping from the complete image collection to these 699 observations remains unspecified. A separate 20-case application table is analyzed only as a supplementary check.

Figure 1 distinguishes the image-processing pipeline from the recorded evaluation sequence. The figure is a reconstructed methodological schematic; it does not imply that every interface callback implements the complete trained-network path.

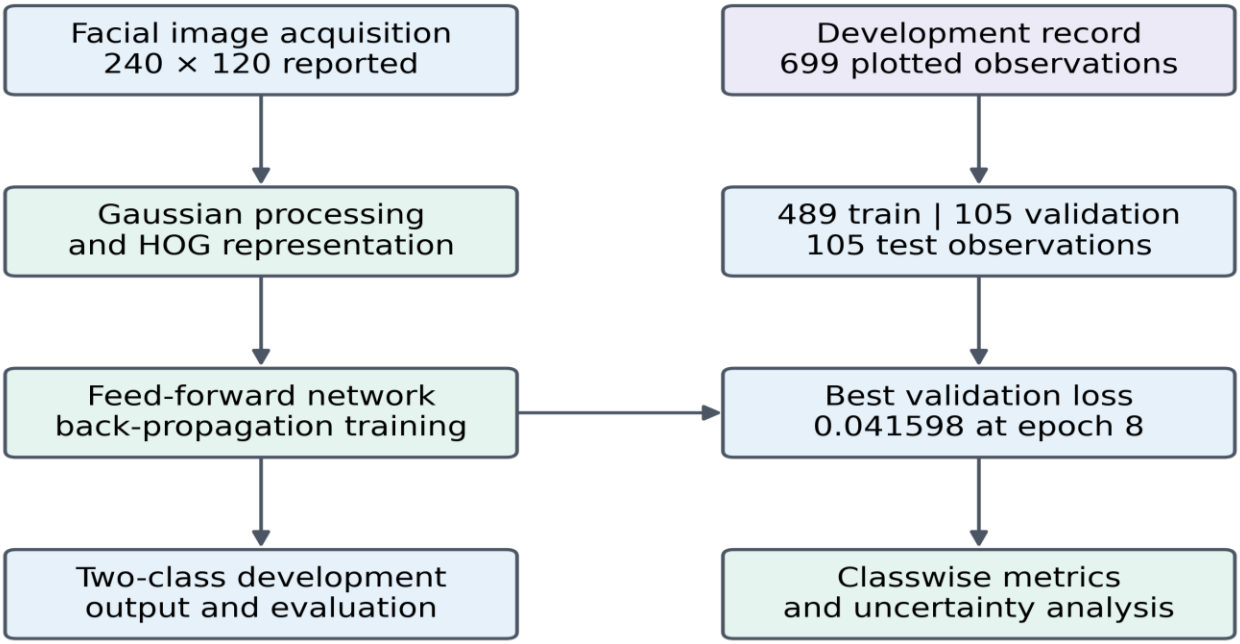


Figure 1. Processing and evaluation sequence reconstructed from the study. The 699 plotted observations are distinct from the 4,030-image collection description.

Table 2: *Development partitions reconstructed from the original confusion figure*

Partition	Target Class 1	Target Class 2	Total	Share of 699
Training	319	170	489	69.96%
Validation	72	33	105	15.02%
Test	67	38	105	15.02%

Note. Class names are retained as numbered in the original figure; no pensioner or impostor meaning is assigned.

3.2 Image processing and gradient features

Let I(u,v) be the intensity at pixel coordinates u and v. A Gaussian kernel of scale σ is defined in Equation (1). The processed image J is the convolution in Equation (2), evaluated over a finite support Ω in implementation.

$$G_\sigma(a,b) = \frac{1}{2\pi\sigma^2}\exp\left[-\frac{a^2+b^2}{2\sigma^2}\right] \quad (1)$$

$$J(u,v) = \sum_{(a,b)\in\Omega} G_\sigma(a,b)\, I(u-a, v-b) \quad (2)$$

Here a and b are kernel offsets, σ controls smoothing, and Ω is the kernel support. The recorded study specifies Gaussian processing without providing the kernel scale or support. These quantities remain model parameters rather than assumed experimental settings. The gradient magnitude m and angle θ follow Equation (3).

$$m(u,v) = \sqrt{g_x{}^2 + g_y{}^2}, \quad \theta(u,v) = \operatorname{atan2}(g_y, g_x) \quad (3)$$

The horizontal and vertical derivatives are g_x and g_γ. For cell c and orientation bin k, Equation (4) accumulates gradient magnitudes with interpolation weight w_k. Neighboring cell histograms form block vector v, normalized in Equation (5).

Corresponding Author Tel: +234 80 6707 3213

$$h_c(k) = \sum_{(u,v)\in c} m(u,v)\, w_k[\theta(u,v)] \quad (4)$$

$$\tilde{v} = \frac{v}{\sqrt{\|v\|_2^2+\varepsilon^2}}, \quad x = \mathrm{concat}(\tilde{v}_1, \ldots, \tilde{v}_B) \quad (5)$$

The small constant ε avoids division by zero; B is the number of blocks, and x is the final descriptor. Its dimension d depends on image geometry, orientation bins, cell dimensions and block overlap. The appendix includes an illustrative HOG extraction call with 32 × 32 cells for a query image. It does not establish the complete trained-model feature configuration. A reproducible implementation should preserve these settings with the model artifact.

3.3 Neural computation and learning objective

The architectural drawing contains two hidden layers using hyperbolic tangent activation. Figure 2 expresses that structure without assigning undocumented layer widths. With $h^{(0)} = x$, the hidden activations are defined by Equation (6).

$$h^{(l)} = \tanh(W^{(l)}h^{(l-1)} + b^{(l)}), \quad l = 1,2 \quad (6)$$

$W^{(l)}$ and $b^{(l)}$ are the weight matrix and bias vector for layer l. The output scores z are converted into two class probabilities by Equation (7). This softmax expression is a conventional probabilistic formulation consistent with cross-entropy classification; the saved MATLAB network is required to confirm its exact output configuration.

$$z = W^{(3)}h^{(2)} + b^{(3)}, \quad p_k = \frac{\exp(z_k)}{\sum_{j=1}^{2}\exp(z_j)} \quad (7)$$

$$L = -\frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{2} y_{ik}\log(p_{ik}) \quad (8)$$

In Equation (8), N is the training sample count and y_{ik} is the one-hot target. Back-propagation computes the gradients of L; a generic parameter update is Equation (9), where Θ collects weights and biases and η is the learning rate. The exact solver and learning rate were not supplied.

$$\Theta_{t+1} = \Theta_t - \eta\,\nabla_\Theta L(\Theta_t) \quad (9)$$

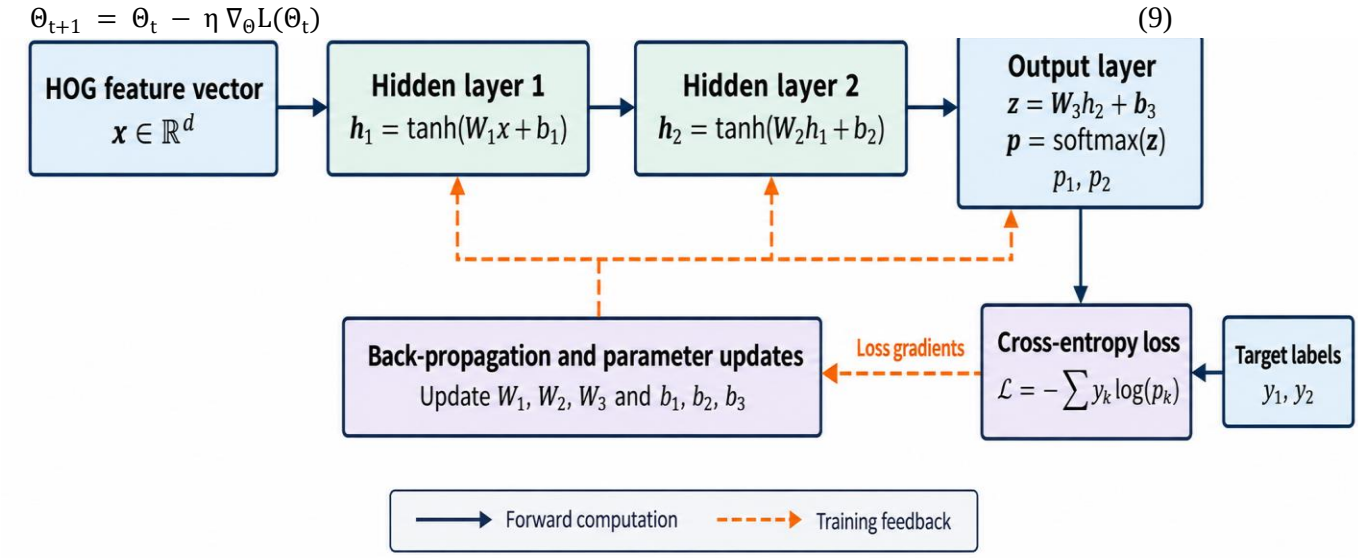


Figure 2. Mathematical organization of the two-hidden-layer model. Layer widths are symbolic; the drawing does not assert undocumented neuron counts.

3.4 Evaluation definitions

The matrices use output class as rows and target class as columns. For analytical convenience only, Class 2 is the positive class: TP = C_{22}, FP = C_{21}, FN = C_{12} and TN = C_{11}. Equations (10) and (11) define the reported measures. This convention does not assign a real-world meaning to either class.

$$A = \frac{TP+TN}{N}, \quad P = \frac{TP}{TP+FP}, \quad R = \frac{TP}{TP+FN} \quad (10)$$

$$F1 = \frac{2TP}{2TP+FP+FN}, \quad BA = \frac{1}{2}\left[\frac{TP}{TP+FN} + \frac{TN}{TN+FP}\right] \quad (11)$$

BA denotes balanced accuracy, which weights the two target-class recalls equally. Count transcription was checked by confirming that the three partition matrices sum exactly to the combined matrix. No score-based receiver-

Corresponding Author Tel: +234 80 6707 3213

operating-characteristic curve or area under that curve is estimated because probability scores are unavailable. Confidence intervals are addressed in the companion reliability analysis.

4 Results and Discussion

4.1 Partition-level classification

Figure 3 redraws the four original matrices at publication resolution. Training yielded 474 correct classifications from 489 observations; validation yielded 101 from 105; and testing yielded 102 from 105. The combined matrix contains 677 correct classifications and 22 errors out of 699 observations. Its 96.85% accuracy rounds to the original 96.9% headline. The test accuracy is 97.14%, rather than a result calculated over all collected images.

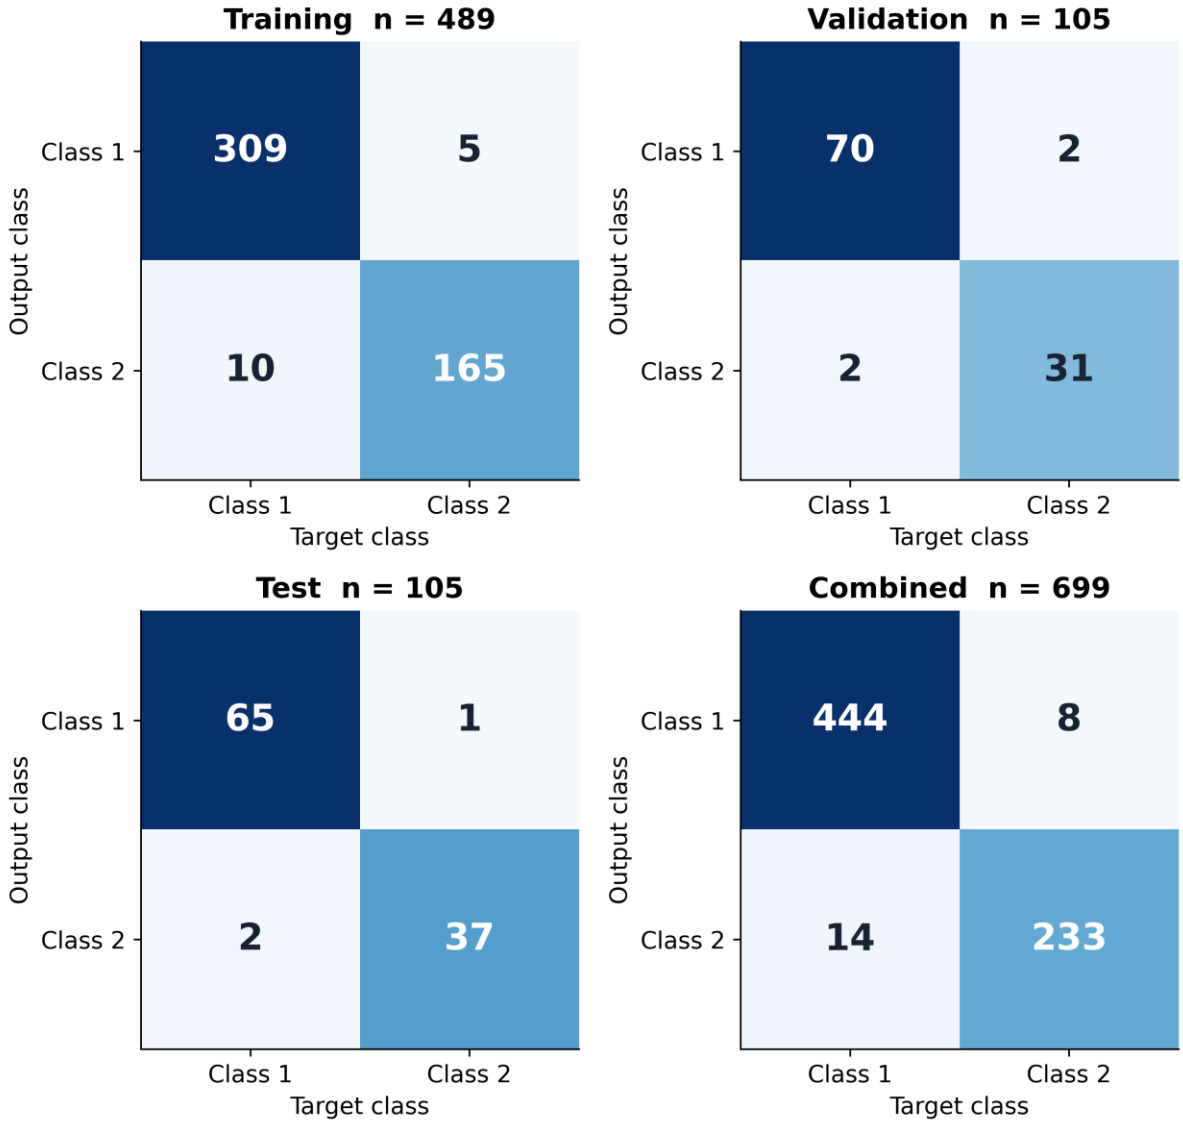


Figure 3. Confusion matrices redrawn from the original study. Rows are outputs and columns are targets. The combined matrix is the arithmetic sum of the three partitions.

Table 3: *Classification measures calculated from the reconstructed matrices*

Partition	Accuracy %	Precision₂ %	Recall₂ %	F1₂ %	BA %
Training	96.93	94.29	97.06	95.65	96.96
Validation	96.19	93.94	93.94	93.94	95.58

Corresponding Author Tel: +234 80 6707 3213

Partition	Accuracy %	Precision$_2$ %	Recall$_2$ %	F1$_2$ %	BA %
Test	97.14	94.87	97.37	96.10	97.19
Combined	96.85	94.33	96.68	95.49	96.81

Note. Subscript 2 identifies the analytical positive class. Combined performance includes training data and is not an independent generalization estimate.

Table 3 shows a narrow descriptive accuracy spread across partitions: 96.19% to 97.14%. On the test partition, the model misclassified two of 67 Class 1 targets and one of 38 Class 2 targets. The resulting Class 2 precision of 94.87% is lower than recall of 97.37%, because the predicted-Class-2 row contains two incorrect assignments but misses one target-Class-2 observation. This is a property of the observed class counts, not evidence that the model admits a particular rate of pension fraud.

Figure 4 separates the four principal measures so that a favorable overall accuracy does not conceal their different denominators. Balanced test accuracy of 97.19% is close to ordinary accuracy, indicating that the observed result is not explained solely by always predicting the more frequent class. An always-Class-1 rule would be correct for 67/105, or 63.81%, of the test observations; that arithmetic reference is not a separately trained baseline.

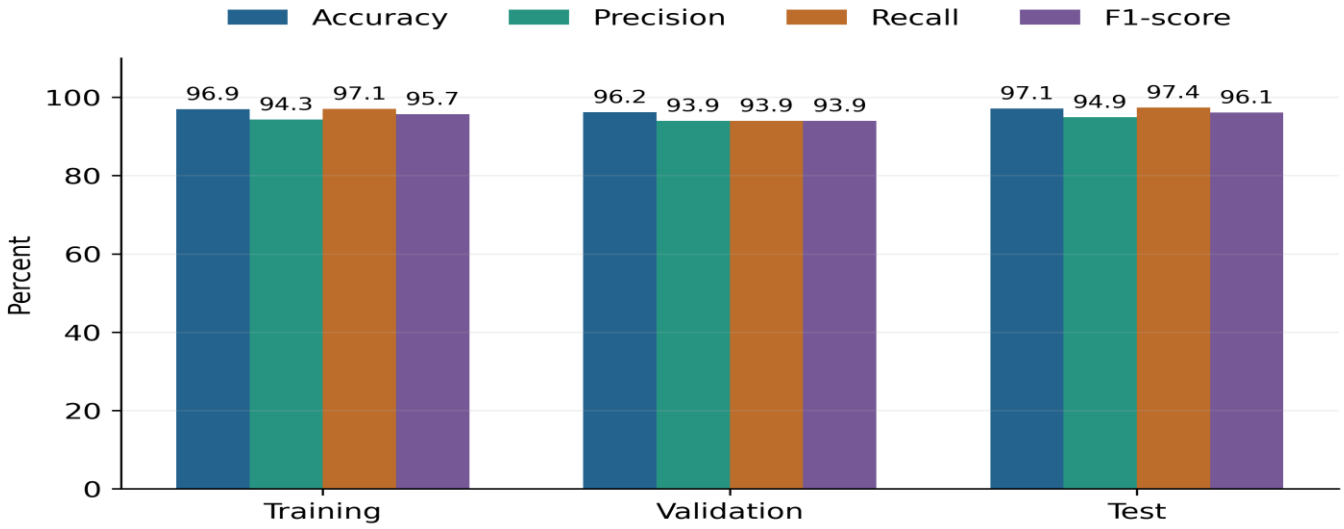


Figure 4. Partition-level accuracy and Class 2 precision, recall and F1-score, calculated from the observed counts. These bars describe one recorded model run.

4.2 Learning behavior

The original learning curve in Figure 5 records a best validation cross-entropy of 0.041598 at epoch 8. The window heading states that training stopped at epoch 14 with a validation-stop condition. Training loss continued downward after the marked validation minimum while the validation trajectory ceased improving. This behavior supports selecting the validation-best checkpoint rather than assuming that the last epoch is best. The six-epoch separation between the marked optimum and stopping point is visible in the record; the configured stopping patience is not reconstructed from that separation.

The curve is reproduced from the study instead of being redrawn from invented intermediate loss values. Cross-entropy reflects probability assignments and differs from the fraction of wrong classifications. Consequently, the 3.1% rounded classification-error complement is not described as a cross-entropy loss.

Corresponding Author Tel: +234 80 6707 3213

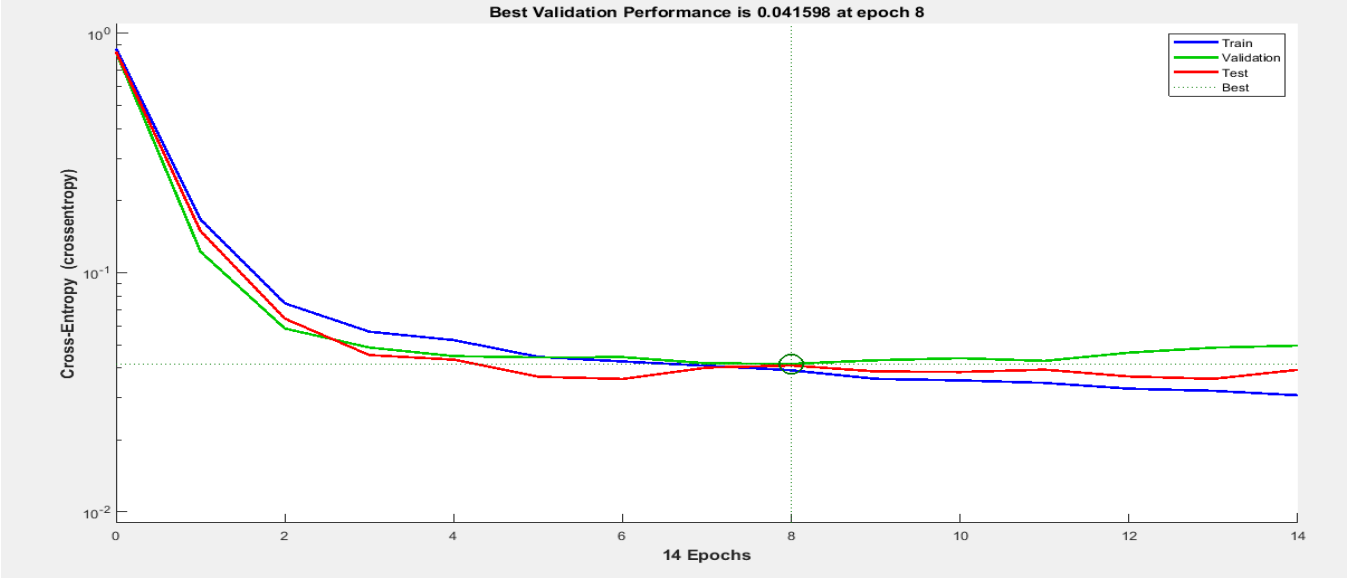


Figure 5. Original training, validation and test cross-entropy plot, cropped to the analytical panel. The marked validation minimum is at epoch 8; the recorded run ends at epoch 14.

4.3 Application relevance and scope

The separate 20-case table contains 18 True and two False recognition entries. Its 90.0% True-coded proportion is retained as an application-level observation. It is not pooled with the development partitions because sample overlap, identity composition and outcome semantics are not established. One False-coded case is described as an impersonator and another as a verified pensioner, so the two entries cannot be converted into operational false-acceptance or false-rejection counts without adjudication.

The principal practical result is evidence of two-class discrimination within the documented development experiment. A pension-management system additionally needs a trusted enrollment link, a claimed identity, a calibrated rejection decision and a review route. Face matching alone also does not demonstrate resistance to presentation attacks (Sharma & Selwal, 2023). Protecting facial templates and limiting record access are separate design responsibilities, as emphasized by Sun and Liu (2025).

The present evaluation comes from one recorded run and lacks session-level partition identifiers, complete hyperparameters and a trained-model export. It does not establish robustness to age change, illumination, occlusion or demographic differences. A follow-up comparison should evaluate the HOG-based network and a modern embedding on the same fixed enrollment and probe records. Repeated seeds and subject-aware confidence estimates would then show how much of the observed performance persists beyond this development record.

5 Conclusion

The study formalized a Gaussian–HOG–feed-forward-neural pipeline and recovered its partition-level classification results. The recorded test partition achieved 102 correct classifications from 105 observations, giving 97.14% accuracy, 96.10% Class 2 F1-score and 97.19% balanced accuracy. The best validation cross-entropy was 0.041598 at epoch 8, and the recorded run stopped at epoch 14. These findings support effective local two-class discrimination and provide a more informative evaluation than the rounded aggregate accuracy alone. Their use in pensioner verification requires confirmed class definitions, traceable enrollment records and independent claim-level testing.

Declarations

Data and software availability. The analysis uses aggregate counts and software excerpts from the underlying study. Identifiable facial images and pension records remain controlled by the original data custodian. Reconstructed count tables and analytical expressions are reported in the article.

Ethics and consent. Institutional ethics approval details, permission to use the pensioner records, and participant-

Corresponding Author Tel: +234 80 6707 3213

consent documentation require confirmation by the original investigators before submission. No identifiable participant photographs are reproduced here.

Funding and competing interests. The authors must confirm the funding statement and any competing interests.

Author contributions. The author list, affiliations, corresponding author and individual research roles require confirmation by the research team.

Use of generative AI. Generative AI assisted language editing, analytical scripting and figure preparation. The authors are responsible for verifying the reported evidence, calculations and final manuscript.

References

Dalal, N., & Triggs, B. (2005). Histograms of oriented gradients for human detection. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Vol. 1, pp. 886–893). IEEE. https://doi.org/10.1109/CVPR.2005.177

Deng, J., Guo, J., Xue, N., & Zafeiriou, S. (2019). ArcFace: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 4690–4699). https://openaccess.thecvf.com/content_CVPR_2019/html/Deng_ArcFace_Additive_Angular_Margin_Loss_for_Deep_Face_Recognition_CVPR_2019_paper.html

Getreuer, P. (2013). A survey of Gaussian convolution algorithms. *Image Processing On Line, 3*, 286–310. https://doi.org/10.5201/ipol.2013.87

Grother, P., Ngan, M., & Hanaoka, K. (2019). *Face recognition vendor test part 3: Demographic effects* (NISTIR 8280). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.8280

Jain, A. K., Ross, A., & Prabhakar, S. (2004). An introduction to biometric recognition. *IEEE Transactions on Circuits and Systems for Video Technology, 14*(1), 4–20. https://doi.org/10.1109/TCSVT.2003.818349

Kim, M., Jain, A. K., & Liu, X. (2022). AdaFace: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 18750–18759). https://openaccess.thecvf.com/content/CVPR2022/html/Kim_AdaFace_Quality_Adaptive_Margin_for_Face_Recognition_CVPR_2022_paper.html

Meng, Q., Zhao, S., Huang, Z., & Zhou, F. (2021). MagFace: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 14225–14234). https://openaccess.thecvf.com/content/CVPR2021/html/Meng_MagFace_A_Universal_Representation_for_Face_Recognition_and_Quality_Assessment_CVPR_2021_paper.html

Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 815–823). https://openaccess.thecvf.com/content_cvpr_2015/html/Schroff_FaceNet_A_Unified_2015_CVPR_paper.html

Sharma, D., & Selwal, A. (2023). A survey on face presentation attack detection mechanisms: Hitherto and future perspectives. *Multimedia Systems, 29*, 1527–1577. https://doi.org/10.1007/s00530-023-01070-5

Sun, Z., & Liu, Z. (2025). Ensuring privacy in face recognition: A survey on data generation, inference and storage. *Discover Applied Sciences, 7*, Article 441. https://doi.org/10.1007/s42452-025-06987-2

Wang, M., & Deng, W. (2021). Deep face recognition: A survey. *Neurocomputing, 429*, 215–244. https://doi.org/10.1016/j.neucom.2020.10.081

Corresponding Author Tel: +234 80 6707 3213