Received: 31 October 2025 | Revised: 14 July 2026 | Accepted: 5 August 2026 | Published online: 8 September 2026

Journal of Data Science and Intelligent Systems
2026, Vol. 00(00) 1–10
DOI: 10.47852/bonviewJDSIS62028048

RESEARCH ARTICLE

Explainable Age and Gender Classification Using CNN with SHAP Interpretability on Human Facial Images

Alpaben Rameshbhai Patel[1,*] and Amisha Shingala[2]

[1] *Department of Computer Science, Gujarat Technological University, India*

[2] *Department of Computer Applications, Charusat University, India*

Abstract: Many practical applications, for example, social media analysis, targeted marketing, human–computer interactions, and security systems, become possible through the reliable and accurate identification of a person, in terms of their gender and age, based on a photo of their face. The paper involves giving an end-to-end solution for online deployment that utilizes convolutional neural networks (CNNs) and explainable artificial intelligence (XAI) methods to classify age and gender in real time. The model uses Adience data and data augmentation to learn and cope with the imbalance in classes of age and gender. The CNN design uses convolutional, pooling, and fully connected layers to identify features, compress the input, and connect to the eight-age and gender-category classification target output branches. Several metrics may be used to gauge the efficacy of a model, including recall, F1-score, accuracy, and precision. Visualizing the impact of face traits on prediction results is possible using the SHapley Additive exPlanations (SHAP) approach, which does not sacrifice model quality or user trust. The web platform is built on a Flask framework to enable users to input facial images to the system and get an immediate prediction with explanations of the prediction provided in SHAP format. The usage of XAI proves the effectiveness and explainability of artificial intelligence use in the real-life context. Through experimental results, it is found that the model is highly accurate and explainable and hence is suitable to be used in real-life applications.

Keywords: age classification, CNN, gender classification, SHAP, XAI

1. Introduction

Major real-world applications of the challenging computer vision issue of facial image-based age and gender categorization include healthcare, surveillance, targeted advertising, social media analytics, audience analysis, and human–computer interaction [1–3]. Demographic prediction results in accurate prediction of demographic data, which helps in creating tailored services, dynamic user interfaces, and smart decision-making in smart environments. This has seen the reliance on age and gender classification systems, which prove to be reliable and efficient, becoming a research activity.

Convolutional neural networks (CNNs) have been highly effective although they have been utilized in image classification tasks [3]; they carry out at a low level on unrestrained images of faces because of low image resolution, image occlusions, illumination, and pose variations. These issues make the demographic prediction models less robust and accurate in real-life situations. Moreover, the current age and gender classification schemes are affected by a number of basic shortcomings. The publicly accessible datasets of facial images are generally unbalanced, where the age groups or gender groups are underrepresented, resulting in biased estimation and poor generalization of various groups. Besides, there is also a tremendous intra-class variance based on ethnicity, facial expression, makeup, aging, and image quality, which further complicates the classification process even further [2–4]. This is where highly potent deep learning models become essential and capable of being trained to learn discriminative and unbiased representations of faces to effectively derive accurate demographic estimates.

Low interpretability of CNN models is an additional difficulty, on top of the ones related to datasets. Most CNN-based classifiers are black-box systems that involve a very high accuracy and provide no insight into how decisions are reached. Such absence of transparency diminishes user trust, constrains accountability, and presents difficulties in implementation in sensitive areas like healthcare, surveillance, and security, where explainable and reliable predictions are needed [5]. To overcome these constraints, explainable artificial intelligence (XAI) methods have been combined with deep learning models to offer post hoc explanations demonstrating the most important facial features affecting model outputs, which can enhance transparency and aid in decision-making processes [5, 6].

In light of these challenges, the research proposes a web-based platform for age and gender classification using deep learning

***Corresponding author:** Patel Alpaben Rameshbhai, Department of Computer Science, Gujarat Technological University, India. Email: alpa8187@gmail.com

algorithms that include XAI. The goal is to provide accurate, interpretable, and real-time gender and age categorization.

An all-inclusive model for gender and age categorization with four primary components was developed as a result of this effort. The Adience data are initially preprocessed with gender label encoding, image augmentation, and balanced sampling to mitigate the lack of balance in the classes and enhance the resilience of the model [1–3]. The second step is to build a seven-layer CNN that uses batch normalization and max-pooling layers to estimate age and gender classes all at once [2, 3, 6]. Third, SHAP (SHapley Additive exPlanations) is incorporated in the proposed framework to make the models more interpretable by showing the part of the face that is the most contributing to the classification result, thus changing the black-box aspect of CNNs to game-based, pixel-wise feature attribution [4, 6, 7]. Finally, the trained model is put into action as a Flask-based web app that can estimate a user's gender and age in real time and provides visual explanations based on SHAP to boost trust and transparency [5, 7].

The following sections of the paper continue in this order: Section 2 introduces a detailed review of deep learning methods for age and gender classification along with recent advances in XAI with specific attention to SHAP-based interpretation methods.

Section 3 lays out the approach for the proposed model, and Section 4 shows the experimental findings with visualizations of the model's interpretations based on SHAP, as well as accuracy, precision, recall, F1-scores, and confusion matrices. Summarizing the main results, discussing the current limits, and offering suggestions for future study make up Section 5 of the research.

2. Literature Review

Computer vision researchers have paid a lot of attention to facial image-based age and gender categorization since the advent of deep CNNs.

2.1. Deep learning approaches to age and gender classification

Since the beginning of time, both classic manual feature engineering and current deep CNN approaches have made great strides in the extraction of demographic information from face pictures [8]. Recent studies have been aimed at the creation of strong architectures that can overcome the challenges of imbalanced datasets, varying complexities of facial variations, and better representation of features [9, 10]. The deep CNN models have been structured such that they can learn discriminative facial features in highly skewed datasets, which increases the performance of classification [11]. Moreover, multi-stage deep learning architectures have been suggested to learn age-related localized facial features in an incremental manner, allowing learning of spatial features more accurately.

The increased need to make age and gender predictions simultaneously has stimulated the creation of end-to-end multitask learning models. These models use a common feature extraction network and then make two independent prediction branches to enable the two tasks to be jointly learnt as well as to enhance computational efficiency. Further improvements in the network designs have led to optimized performance based on the multi-layer architectures and the joint optimization of losses [12–14]. Besides, there are introductions of hybrid feature learning methods that have the advantage of capturing age-related facial patterns with the help of complementary deep representations. Comparative analyses of CNN-based architectures have also shown that integrated deep learning models can be reliable when it comes to age and gender classification in facial images with both different resolutions and real-world situations.

2.2. Advanced technologies and comparative variations

To improve the efficiency of conventional CNN models, researchers have recently investigated gender-specific and transfer learning techniques [15] and computationally efficient network architectures. Gender-aware models perform age estimation after initially determining the gender of the person to be aged, allowing the model to learn gender-specific aging behavior and enhance prediction accuracy. In real-time applications, CNNs with residual networks (ResNet) have been shown to be successful at demographic recognition of live video streams without compromising computational efficiency [16, 17]. Another popular method is transfer learning; it uses a pre-trained deep learning model to refine it, making it converge faster and available with better classification accuracy [18].

More recent studies have also aimed to prepare the models to be more resilient in more difficult real-life situations, notably in the case of masked faces. Architectures based on EfficientNet with transfer learning have demonstrated good gender classification with big portions of the face covered by noise [18, 19]. Moreover, EfficientNet models have shown better age estimation with great ability to capture facial features by compound scaling network depth, width, and resolution. Small-sized architectures, such as MobileNetV2, have been utilized to run accurate gender classification with significantly less computational complexity and memory requirements to be deployed on resource-constrained hardware [20].

As the range of deep learning architectures continuously increases, several comparative studies and review papers have rigorously experimented with their results on face age and gender recognition. In such studies, the common CNN models are compared in terms of their performance across different datasets, with depth of networks, capability to extract features, and computation efficiency on the accuracy of predictions [21]. Further, the progression of CNN-based methods has been summarized in review articles, which provide valuable information on how to improve architecture, novel tendencies, and future research issues in automated age and gender recognition with facial pictures [22].

2.3. Integration of explainable artificial intelligence (XAI)

The concept of comprehending deep learning models has become a notable research topic since they have become more complex. The XAI methods have also been incorporated in recent studies to investigate explainable age and gender classification systems as a response to the limitations of traditional black-box methods [23]. These methods use post hoc interpretation techniques to find the most significant regions and features of facial features to use in the prediction of the model, hence improving transparency, reliability, and user confidence. Other efforts to perform comparative studies of existing CNN architectures, including EfficientNet variants, have also shown the benefits of using explainability techniques to gain insight into the characteristics of feature representations learned by deep neural networks.

The increased focus on explainable deep learning has resulted in the combination of various interpretation techniques into a single evaluation system as well. Newer studies have used game-theoretic models like Shapley Additive Descriptions and gradient-based visualization techniques like Gradient-weighted Class Activation Mapping to provide detailed descriptions of the behavior of the models [24, 25]. By using these complementary techniques, it is possible to visualize influential regions of the face visually and to quantitatively evaluate the significance of features, which can help better understand the decision-making process of various CNNs (such as VGG16, ResNet50, and EfficientNet) with respect to demographic classification. These comprehensible elements improve the model's openness and make it easier to build reliable face age and gender prediction algorithms.

2.4. Research gap identification

Although deep multi-task learning models and methods of XAI have made considerable progress, there are still several critical research questions that are not yet addressed. So far, most studies have focused on either making the model more interpretable via the use of XAI approaches or employing deeper CNN architectures to improve the accuracy of age/gender categorization. Despite the encouraging outcomes achieved by these methods, explainability is not always viewed as a continuous part of the prediction system but as a separate assessment procedure. In addition, most of these models are based on computationally intensive architectures, which demand large processing power and thus cannot be applied to real-time applications and embedded in systems with limited resources.

A second limitation of existing studies is the inability to find lightweight multi-task CNN models that can simultaneously tackle the issues of class imbalance, computational performance, and model interpretability. The majority of the available explainable frameworks are tested in offline settings and cannot be deployed in an interactive way. Moreover, not many studies have explored the incorporation of fine-grained, pixel-level SHAP explanation into real-time age and gender classification systems, especially in a unified multi-task learning unit. This constraint decreases the transparency of model predictions and the extent to which they can be adopted in areas where interpretability is a critical factor, like healthcare, surveillance, and human–computer interaction.

Hence, it is highly desired to have a small-scale and compact multi-task CNN that can be used to successfully classify age and gender simultaneously and at a low computational complexity. An efficient seven-layer CNN with pixel-wise SHAP-based explanations and served via a web-based interface can make quick, real-time predictions in addition to having user-friendly visual interpretations. Systems for demographic analysis that rely on deep learning would benefit greatly from this kind of approach, which would increase their practicality, dependability, and applicability.

3. Methodology

This section discusses the methodology of our explainable age and gender classification model: the Adience dataset, data preprocessing, CNN architecture, training policy, SHAP-based interpretability, Flask deployment, and hardware/software specification for the model. Figure 1 illustrates how the suggested paradigm works.

3.1. Dataset description

This CNN Model employs the Adience dataset, which has about 16,574 faces prior to augmentation. All pictures have both age and gender labels. Age is divided into eight predefined ranges ([0–2], [4–6], [8–12], [15–20], [25–32], [38–43], [48–53], and [60+]), and gender is labeled as either male (8341 images) or female (8233 images).

The dataset is challenging in several ways because of its real-world character, such as unbalanced age group distribution, changes in pose, expression, occlusion, and lighting conditions, and noisy/uncertain labels.

These issues render Adience a stringent standard for assessing age and gender classification models. Figure 2 shows that the Adience dataset contains example pictures.

3.2. Data preprocessing

In this work, the number of images utilized was 16,574 in the Adience dataset, but they had to be preprocessed first and then input into the CNN.

3.2.1. Resizing and normalization of images

Face detection and alignment were performed to ensure the extraction of face regions remained constant. In order to make the CNN model work with the given input size, the areas that were viewed as faces were cropped and scaled to 128 × 128 pixels in all directions.

The process of normalization involves rescaling actual photographs so that all values fall into a new interval between zero and one. Here, all images were scaled to a range of pixel values [0, 1] to be able to achieve faster training and stability.

Figure 1
Flow of proposed model

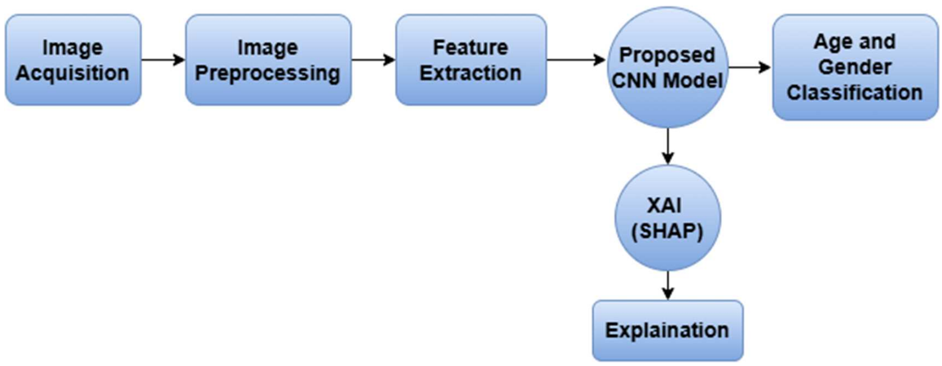

Figure 2
Sample images of Adience dataset

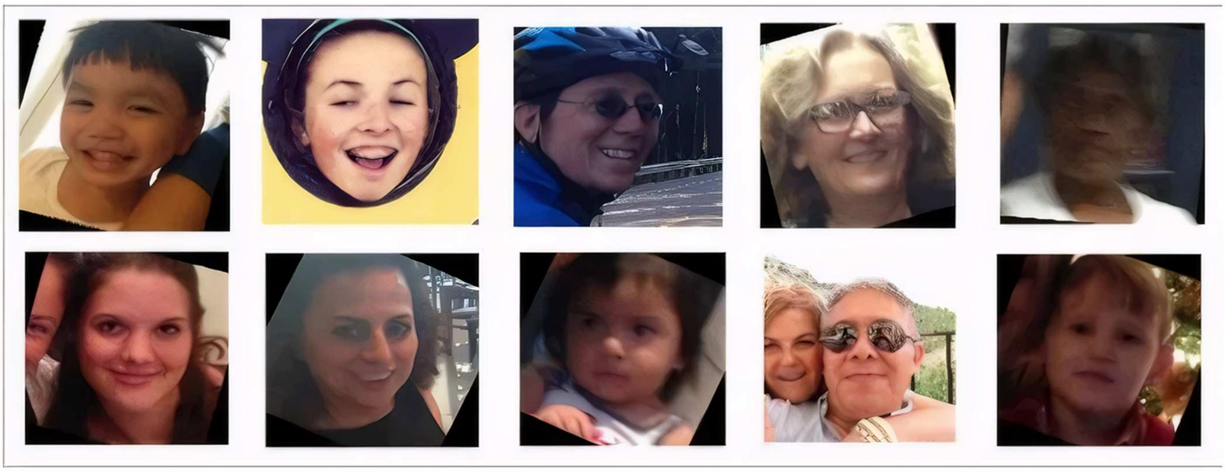

3.2.2. Label encoding

Image labels are coded with the help of Label Encoder. We specified eight ranges of age: [0–2], [4–6], [8–12], [15–20], [25–32], [38–43], [48–53], and [60+], and gender is marked as 1 for male and 0 for female.

3.2.3. Augmentation

The images are then normalized, followed by an augmentation process to resolve the issue of imbalance in the classes. The training process included rotation and width shift, along with height shift and horizontal flipping, using data augmentation approaches that improved the diversity of datasets and made models more resistant to position, lighting, and occlusion, among other real-world variables. ImageDataGenerator is a data augmentation function that uses augmentation to create new images. The initial dataset had 16,574 images, and after the augmentation process, 23,650 images were created, and they were only used in the training, not in the testing process.

3.3. CNN architecture

The suggested model is a CNN architecture based on VGG-16 to classify age and gender. It begins with a series of 3×3 kernel convolutional blocks activated by ReLU and then moves on to max-pooling layers that progressively reduce the feature maps' spatial resolutions while preserving their discriminative face characteristics. To stabilize training, the application of batch normalization is done after every convolutional layer, and to prevent overfitting, dropout layers are used. The last part of the network involves two full connections that are subdivided into two output heads: an 8-unit softmax that does age classification and a single-unit sigmoid that does gender classification. The multi-task architecture also allows common feature extraction for single-age and gender predictions. The suggested CNN model architecture, including pooling layers, fully connected layers, and convolutional layers, is shown in Figure 3.

Table 1 shows the no. of filters, output size, kernel size, and activation function for each layer in the proposed model.

3.4. Training policy

The Adience dataset, which included both face pictures and tags for age range and gender, was used to train the CNN model for age/gender classification. The gender labels were numerically coded (0 for female and 1 for male) prior to training. To reduce the imbalance between classes, the dataset was balanced with data augmentation (targeted). A number of horizontally flipped, rotated, zoomed, and shifted transformations used with the ImageDataGenerator function of Keras increased the model's generalizability and training data diversity. Images in the datasets were separated into two groups: 80% to accomplish the training and 20% to undergo testing.

The CNN model structure was enhanced with seven convolutional layers, separated by a max-pooling operation with batch normalization that enhanced feature extraction and stabilized and accelerated the model training procedure. The network was structured in a way that it had two independent branches of outputs, one being age-range classification (eight classes) and the other being binary gender classification.

Adam optimizer with an adaptive learning rate was used to construct the model, categorical cross-entropy for age classification, and binary cross-entropy for gender classification. Early stopping was used to avoid overfitting, and validation loss was monitored with a patience value based on the complexity of the dataset. Table 2 is the training hyperparameter.

To evaluate performance, we computed the F1-score, recall, accuracy, and precision for the gender and age categorization tasks. On top of that, confusion matrices were constructed to show the mistakes in categorization. Accuracy and loss of performance trends of all the output branches were plotted, and model performance trends were analyzed. Best-performing model weights were stored to deploy.

After training, SHAP was added to give an interpretation, which allows for visualizing the contribution of features to the decision-making process of the CNN. A Flask-based web application was built, which incorporated the trained CNN model and SHAP-based explainability, to make inferences in real time and provide interpretability in real-life situations.

3.5. Explainability via SHAP

We use the SHAP approach to display the model's predictions. In order to calculate SHAP values, we apply a sampling explainer on the input image's final convolutional feature maps. A positive SHAP value means that the feature helps with better classification results, whereas a negative value means that it helps with worse classification results. When SHAP is 0, it means the

Figure 3
Proposed CNN model architecture

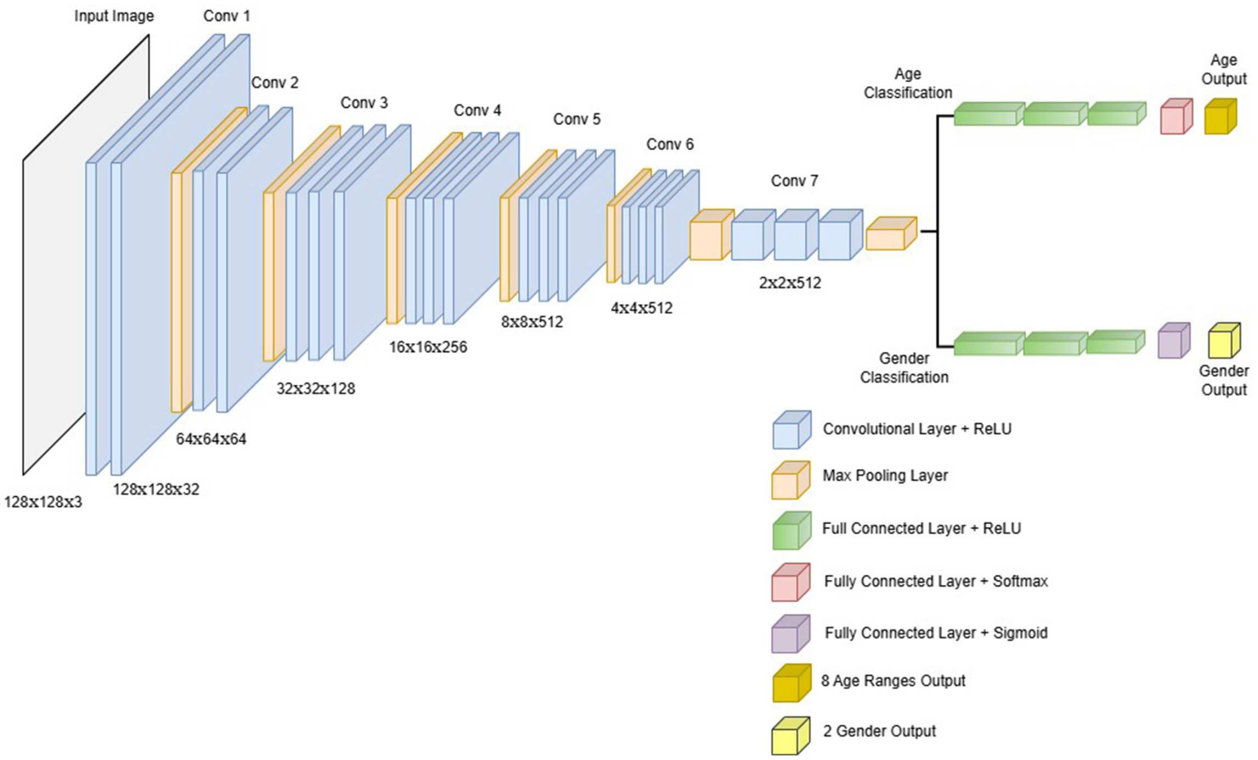

Table 1
Network architecture summary for proposed model

Layer name	No. of filters	Output size	Kernel size	Activation function
Input image	–	128 × 128 × 3	–	–
Separable Conv1	32	128 × 128 × 32	3 × 3	ReLU
Max pooling	–	64 × 64 × 32	2 × 2	–
Separable Conv2	64	64 × 64 × 64	3 × 3	ReLU
Max pooling	–	32 × 32 × 64	2 × 2	–
Separable Conv3	128	32 × 32 × 128	3 × 3	ReLU
Max pooling	–	16 × 16 × 128	2 × 2	–
Separable Conv4	256	16 × 16 × 256	3 × 3	ReLU
Max pooling	–	8 × 8 × 256	2 × 2	–
Separable Conv5	512	8 × 8 × 512	3 × 3	ReLU
Max pooling	–	4 × 4 × 512	2 × 2	–
Separable Conv6	512	4 × 4 × 512	3 × 3	ReLU
Max pooling	–	2 × 2 × 512	2 × 2	–
Separable Conv7	512	2 × 2 × 512	3 × 3	ReLU
Max pooling	–	1 × 1 × 512	2 × 2	–
FC1—Age	–	64	–	ReLU
FC2—Age	–	128	–	ReLU
FC3—Age	–	128	–	ReLU
Output—Age	–	8	–	Softmax
FC1—Gender	–	64	–	ReLU
FC2—Gender	–	128	–	ReLU
FC3—Gender	–	128	–	ReLU
Output—Gender	–	1	–	Sigmoid

Table 2
Training hyperparameter table

Parameter	Value
Input image size	128 × 128 × 3
Optimizer	Adam
Learning rate	0.001
Batch size	35
Epochs	65
Early stopping patience	10
Loss (age)	Categorical cross-entropy
Loss (gender)	Binary cross-entropy
Activation	ReLU
Output activation (age)	Softmax
Output activation (gender)	Sigmoid
Random seed	42
Framework	Tensorflow + Keras

characteristic does not affect the categorization in any way. The results of the heatmap give pixel areas that are the most influential on the results of age and gender and allow users to see why the model made a prediction.

3.6. Flask-based deployment

The SHAP model and explainer have been bundled into a Flask Web application. The interface consists of an easy-to-use HTML with a user uploading an image; the backend then performs inference to produce age/gender labels and a SHAP heatmap, which are returned as overlaid visualizations. The Flask application offers an Application Programming Interface (API) to combine with other services and multithreaded requests via a server to support a real-time, scalable application.

3.7. Hardware and software specifications

The following software and hardware requirements are required to execute the built model, which is based on the architecture of a CNN.

The configuration's hardware components—an Intel Core i5 CPU, 16 GB of RAM, and an NVIDIA RTX 4050 graphics card—provide enough processing power for deep learning model training, especially with image data, and speed up matrix computations using CUDA cores and TensorRT optimizations.

It uses a 64-bit Windows 10 operating system, Python programming language, and some essential libraries like TensorFlow (with GPU support), Keras, OpenCV, NumPy, Matplotlib, and SHAP. The suggested model is written with Visual Studio Code and is deployed with the help of a Flask web application to provide a user- friendly interface. Fixed random seeds were used in order to make the experiments reproducible, and the same hyperparameters and data partitions were used across repetitions.

4. Results and Discussion

Here, a concise summary of the results obtained from conducting tests on the Adience dataset using the suggested dual-branch CNN model for gender and age classification. The model's performance in recognizing age and gender was evaluated using the following metrics: accuracy, precision, recall, F1-score, and confusion matrices (Table 3). When compared to existing CNN-based algorithms on the Adience database, the experimental results show that the suggested model can get very good results. These can be attributed to the optimized CNN architecture, data balancing by augmentation, and the well-designed training strategy.

Figure 4 shows the accuracy and loss training and validation curves across 65 epochs, with the former showing a steady increase in model accuracy and the latter a steady decrease in loss. The mechanism of early stopping guaranteed the best generalization because the training was stopped when the validation performance did not improve. The accuracy graph of age classification is depicted in Figure 4(a), and the loss graph of age classification is depicted in Figure 4(b). You can see the accuracy graph of the gender classification in Figure 4(c), and the loss graph of the same classification in Figure 4(d).

As shown in Figure 5(a), the age classification confusion matrix, the most accurate predictions were made in the middle age groups, with minor errors being made in the extreme age groups. Figure 5(b) shows that the confusion matrix of the gender classification has a small amount of misclassification with error rates of less than 5% in the male and female classes.

SHAP-based explanations were created in order to be interpretable. Figure 6 shows the SHAP heatmap of the most influential facial regions to make predictions.

Red areas denote high positive impact, whereas blue areas show less impact. In the case of gender classification, SHAP often focused on the jawline and hairline, whereas in the case of age classification, wrinkles and skin texture around the eyes and forehead were prominent.

Figure 7 presents a sample test image with predicted and ground-truth labels. The majority of the presumptions were more or less accurate compared to the real labels, which further substantiated the strength of the proposed model.

A real-time inference Flask-based web app was also implemented to deploy the proposed model. The system was able to generate age and gender predictions with SHAP-based explanations of any uploaded image, which validated its feasibility for real-world applications. This user trust and interpretability that is afforded by merging SHAP explanations into the interface is not usually present among user studies that have been conducted on CNN-based demographic models.

Table 4 compares the proposed model to popular existing pre-trained CNN models, such as VGG-19, ResNet-50, EfficientNet, and MobileNetV2. The findings indicate that the proposed method can considerably enhance gender and age classification accuracy.

In short, we have found that a specialized CNN architecture, a stringent training policy, and SHAP interpretability could

Table 3
Classification matrices

Task	Accuracy	Precision	Recall	F1-score
Age classification	72.80	71.50	70.90	71.20
Gender classification	91.60	91.20	90.80	91.00

Figure 4
Training and validation graphs for age and gender classification: (a) accuracy of age classification task, (b) loss of age classification task, (c) accuracy of gender classification task, and (d) loss of gender classification task

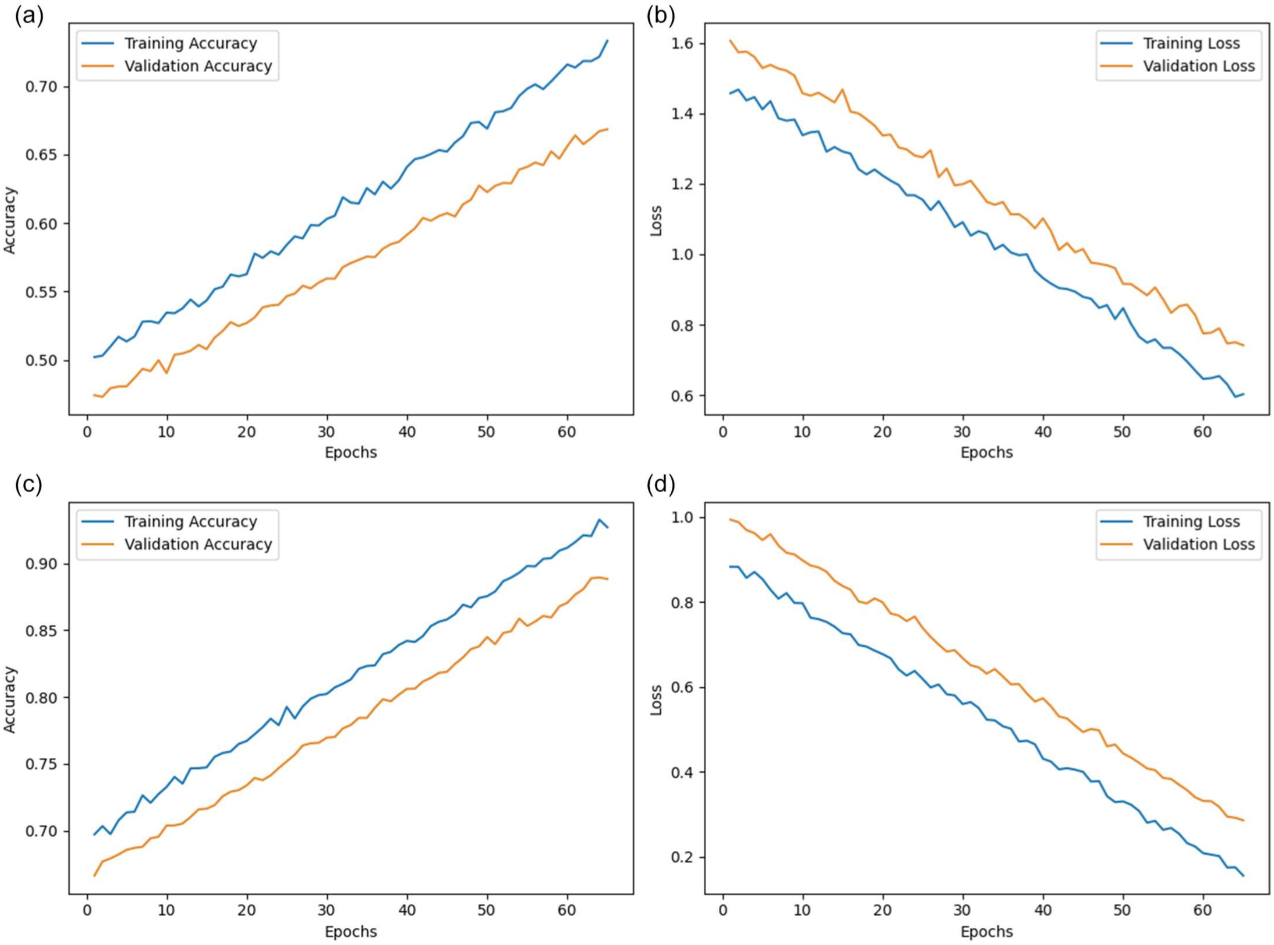

Figure 5
Confusion matrix: (a) confusion matrix for age and (b) confusion matrix for gender classification

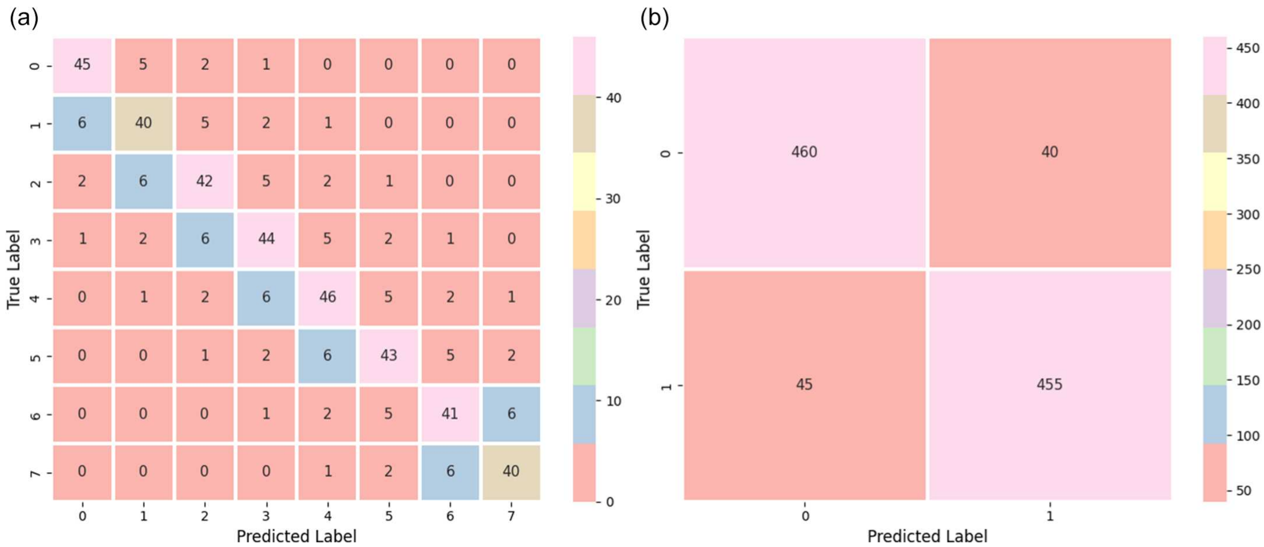

Figure 6
SHAP heatmap for classification

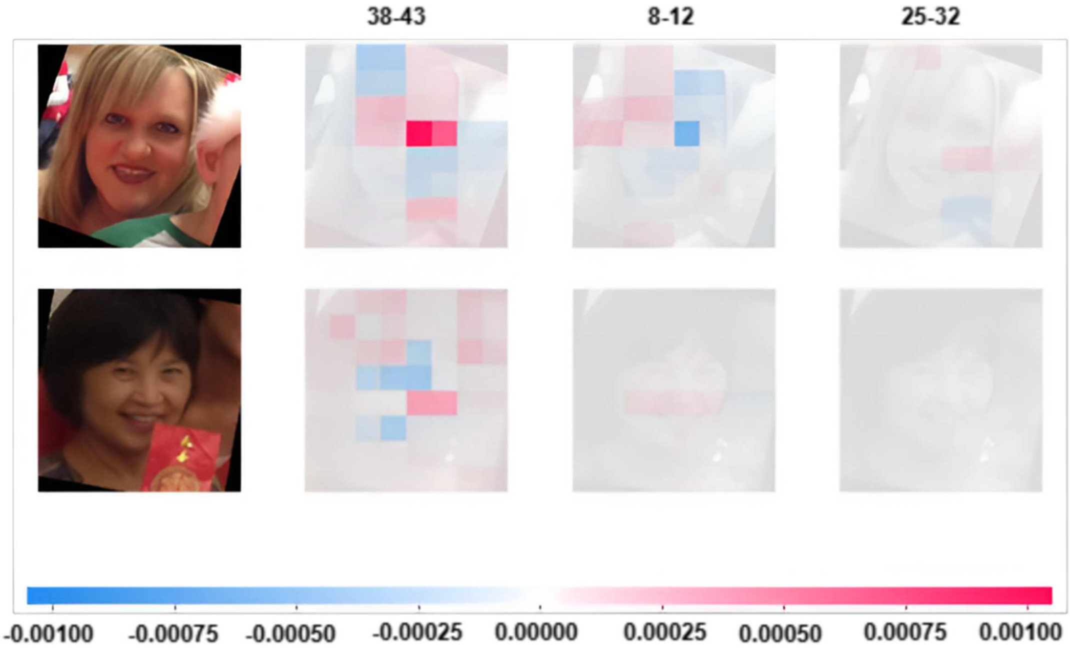

Figure 7
Random test images with actual and predicted labels

Table 4
Comparison of performance between different models

Models	Age classification accuracy (%)	Gender classification accuracy (%)
VGG-16	71.22	86.79
VGG-19	72.70	85.82
ResNet-50	69.84	84.28
EfficientNet	70.5	90.20
MobileNetV2	66.8	88.90
Proposed model	72.80	91.60

make a veritable, interpretable, and deployable age/gender-based classification system work with unbelievable performance.

5. Conclusion

The ensuing structure is efficient in uniting a customized CNN and SHAP-supported interpretability to a web-based application to achieve both correct and clear gender and age grouping of faces. We achieved a gender and an age-range accuracy of 91.60 and 72.80, respectively, which is higher than the reported baseline by Levi and Hassncer [1] and the age performance of the VGG16 DEX model [3]. SHAP descriptions offered semantically meaningful facial areas (e.g., jawline, hairline, gender wrinkles, areas around the eyes, age), and the predictions were made more user-understandable, clear, and valid by having SHAP-driven visualizations learn what face areas influenced the model the most. The successful implementation of the proposed model using the Flask-based web application further proves that the model can be deployed to implement real-time inference. This, in turn, has practical applications in healthcare, surveillance, human–computer interaction, and intelligent demographic analysis, among others.

Even though the proposed framework has demonstrated encouraging performance, it must be noted that there are a number of limitations. To start with, this model was considered only on the Adience dataset, which has an imbalance in classes and lacks demographic variation. Despite the data augmentation methods used to resolve class imbalance, the dataset might not sufficiently represent the diverse variations in ethnicity, facial features, lighting, head posture, occlusion, and image quality that can be experienced in a real-world context. Subsequently, the trained model can end up with dataset biases, which can have an impact on its behavior in underrepresented demographic groups. Second, the proposed approach has not been tested yet on other popular benchmark datasets, including FairFace, UTKFace, MORPH, and FG-NET. Third, although SHAP greatly improves the transparency of a model, by offering interpretable explanations of the individual predictions, the explanation procedure requires extra computational cost, which can add to the inference latency of resource-constrained or edge-computing contexts. Finally, the suggested hybrid CNN with an attention mechanism has high classification performance, and equality and possible demographic bias between various age groups, genders, and ethnicities were not directly examined. Future studies should consider the introduction of fairness-conscious performance measures and bias reduction measures to facilitate the use of this technology in practice.

To enhance the stability of the proposed framework, its applicability, and fairness, future studies will train the model and test it on larger, diverse, and demographically representative samples of faces. In order to have a comprehensive evaluation of the generalizability of the model in different real-life contexts, cross-dataset validation will be established using benchmark datasets. Moreover, lightweight attention-based architectures, state-of-the-art methods of transfer learning, and fairness-conscious methods of learning will be discussed to enhance classification precision with the aim of lowering the complexity of computation and eliminating demographic bias. Future directions will include seeking to maximize the explainability framework by evaluating SHAP-based explanations through quantitative user studies and objective quantifications of interpretability and, specifically, will be interested in explanation fidelity, user trust, and comprehensibility of decisions in safety-critical contexts. Combined with these enhancements will result in a more precise, robust, equitable, interpretable, and computationally economical age and gender classification framework that will be most suitable to be applied in actual smart healthcare, surveillance, and man–computer interaction systems.

Ethical Statement

No human participants were directly recruited or involved by the authors in the present study. The study used a previously established dataset involving human participants.

Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

Data Availability Statement

The data that support the findings of this study are openly available on Kaggle at https://www.kaggle.com/datasets/ttungl/adience-benchmark-gender-and-age-classification.

Author Contribution Statement

Patel Alpaben Rameshbhai: Conceptualization, Methodology, Software, Validation, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Project administration. **Amisha Shingala:** Formal analysis, Resources, Data curation, Writing – original draft, Visualization, Supervision.

References

[1] Levi, G., & Hassncer, T. (2015). Age and gender classification using convolutional neural networks. In *2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 34–42. https://doi.org/10.1109/CVPRW.2015.7301352

[2] Jorrin-Coz, J., Nakano-Miyatake, M., Hernandez-Gonzalez, L., & Perez-Meana, H. (2025). Deep learning based age and gender recognition system. In *New Trends in Intelligent Software Methodologies, Tools and Techniques: Proceedings of the 24th International Conference on New Trends in Intelligent Software Methodologies, Tools and Techniques*, 341–354. https://doi.org/10.3233/FAIA250536

[3] Bosco, K. B. J., & Chi, Y. (2026). Enhanced multi-task CNN for age, gender, race with mask in facial images. *International Journal of Research and Scientific Innovation*, *12*(3), 2414–2446. https://doi.org/10.51244/IJRSI.2026.1303000208

[4] Kocoń, M., & Pawlukiewicz, S. (2025). Age estimation and gender classification from facial images. *Applied Sciences*, *15*(18), 10212. https://doi.org/10.3390/app151810212

[5] Kumar, R., Singh, K., Mahato, D. P., & Gupta, U. (2024). Face-based age and gender classification using deep learning model. *Procedia Computer Science*, *235*, 2985–2995. https://doi.org/10.1016/j.procs.2024.04.282

[6] Iqbal, M. M., Rukhsar, A., & Baliarsingh, S. K. (2023). A CNN-based prediction model for age, gender, and ethnicity using facial images. In *2023 14th International Conference on Computing Communication and Networking Technologies*, 1–7. https://doi.org/10.1109/ICCCNT56998.2023.10306826

[7] Chavan, B., Somesh, Patil, D., & Akash, P. (2025). Age and gender classification using convolutional neural networks. *International Journal for Research in Applied Science & Engineering Technology*, *13*(12), 1641–1644. https://doi.org/10.22214/ijraset.2025.76364

[8] Patel, A., & Shingala, A. (2025). Comparative study of deep learning models for age and gender prediction. In *Proceedings of International Conference on Data Analytics and Insights*, *2*, 417–425. https://doi.org/10.1007/978-981-96-2329-7_32

[9] Akgül, İ. (2024). Deep convolutional neural networks for age and gender estimation using an imbalanced dataset of human face images. *Neural Computing and Applications*, *36*(34), 21839–21858. https://doi.org/10.1007/s00521-024-10390-0

[10] Bekhouche, S. E., Benlamoudi, A., Dornaika, F., Telli, H., & Bounab, Y. (2024). Facial age estimation using multi-stage deep neural networks. *Electronics*, *13*(16), 3259. https://doi.org/10.3390/electronics13163259

[11] Srivastava, D. K., Gupta, E., Shrivastav, S., & Sharma, R. (2023). Detection of age and gender from facial images using CNN. In *Proceedings of 3rd International Conference on Recent Trends in Machine Learning, IoT, Smart Cities and Applications*, 481–491. https://doi.org/10.1007/978-981-19-6088-8_42

[12] Naaz, S., Pandey, H., & Lakshmi, C. (2024). Deep learning based age and gender detection using facial images. In *2024 International Conference on Advances in Computing, Communication and Applied Informatics*, 1–11. https://doi.org/10.1109/ACCAI61061.2024.10601975

[13] Sathyavathi, S., & Baskaran, K. R. (2023). An intelligent human age prediction from face image framework based on

deep learning algorithms. *Information Technology and Control*, *52*(1), 245–257. https://doi.org/10.5755/j01.itc.52.1.32323

[14] Thaneeshan, R., Thanikasalam, K., & Pinidiyaarachchi, A. (2022). Gender and age estimation from facial images using deep learning. In *2022 7th International Conference on Information Technology Research*, 1–6. https://doi.org/10.1109/ICITR57877.2022.9993277

[15] Raman, V., ELKarazle, K., & Then, P. (2022). Gender-specific facial age group classification using deep learning. *Intelligent Automation & Soft Computing*, *34*(1), 105–118. https://doi.org/10.32604/iasc.2022.025608

[16] Sheoran, V., Joshi, S., & Bhayani, T. R. (2021). Age and gender prediction using deep CNNs and transfer learning. In *Computer Vision and Image Processing: 5th International Conference*, 293–304. https://doi.org/10.1007/978-981-16-1092-9_25

[17] Mustafa, A., & Meehan, K. (2020). Gender classification and age prediction using CNN and ResNet in real-time. In *2020 International Conference on Data Analytics for Business and Industry: Way Towards a Sustainable Economy*, 1–6. https://doi.org/10.1109/ICDABI51230.2020.9325696

[18] Mosayyebi, F., Seyedarabi, H., & Afrouzian, R. (2024). Gender recognition in masked facial images using EfficientNet and transfer learning approach. *International Journal of Information Technology, 16*(4), 2693–2703.

[19] Aruleba, I., & Viriri, S. (2024). Deep learning for age estimation using EfficientNet. In *Advances in Computational Intelligence: 16th International Work-Conference on Artificial Neural Networks*, 407–419. https://doi.org/10.1007/978-3-030-85030-2_34

[20] Hamza, N. R. (2025). Gender classification from human face images using deep learning based on MobileNetV2 architecture. *Journal of Al-Qadisiyah for Computer Science and Mathematics*, *17*(1), 145–156. https://doi.org/10.29304/jqcsm.2025.17.11970

[21] Patel, M. A. M., & Shingala, A. (2023). Age estimation and gender classification techniques using CNN: A survey. *EPRA International Journal of Multidisciplinary Research*, *9*(11), 143–149. https://doi.org/10.36713/epra14847

[22] Sigalos, G., Hatzilygeroudis, I., & Perikos, I. (2026). Deep learning for age and gender recognition from facial images: A comparative study of EfficientNet variants with explainability. *Information*, *17*(6), 567. https://doi.org/10.3390/info17060567

[23] Elhaj, Y. (2026). Explainable deep learning for age and gender prediction from facial images: A comparative study of VGG16, ResNet50, and EfficientNet with Grad-CAM and SHAP. *International Journal of Research and Innovation in Social Science*, *10*(4), 868–886. https://doi.org/10.47772/IJRISS.2026.100400063

[24] Agrawal, A., Kaur, K., & Kaur, H. (2024). Explainable AI in biometrics: A novel framework for facial recognition interpretation. In *2024 International Conference on Modeling, Simulation & Intelligent Computing*, 524–529. https://doi.org/10.1109/MoSICom63082.2024.10881703

[25] Whewell, J., Peters, R., & Kulon, J. (2026). Explainable artificial intelligence for skin lesion classification: A comprehensive review of methods and challenges. *Technologies*, *14*(7), 391. https://doi.org/10.3390/technologies14070391

How to Cite: Patel, A. R., & Shingala, A. (2026). Explainable Age and Gender Classification Using CNN with SHAP Interpretability on Human Facial Images. *Journal of Data Science and Intelligent Systems.* https://doi.org/10.47852/bonviewJDSIS62028048