OPEN ACCESS
International Journal of Innovations in Science & Technology

Deep Faces: Advancing Age and Gender Classification using Facial Images with Deep Features

Ammad Noor*, Wakeel Ahmad, Syed M. Adnan Shah
Department of Computer Science, University of Engineering and Technology Taxila, Pakistan.
* Correspondence: Ammad Noor, ammadnoor.ms@gmail.com

Citation | Noor. A, Wakeel. A, Shah. S. M. A, "Deep Faces: Advancing Age and Gender Classification using Facial Images with Deep Features", IJIST, Vol. 6 Issue. 2 pp 808-818, June 2024

Received | June 04, 2024 **Revised |** June 22, 2024 **Accepted |** June 25, 2024 **Published |** June 26, 2024.

In the realm of identity recognition and social interactions, human facial features play a pivotal role. Accurate age estimation and gender classification from facial images have practical implications across various fields, including biometrics, surveillance, and personalized services. This study presents a novel approach that harnesses deep features extracted by the VGG-19 architecture for age and gender prediction, employing a custom convolutional neural network (CNN) for classification. Leveraging the UTKFace dataset, encompassing a diverse collection of facial images with annotated age and gender labels spanning various ages, ethnicities, and gender representations, provides a robust foundation for model training and evaluation. Deep features extracted from the VGG-19 architecture serve as rich representations of facial patterns, enabling our model to discern discriminative cues for age and gender. These deep features are input to CNN model, which is fine-tuned specifically for age and gender classification. The model comprises input layer, Dense layers, incorporating dropout and batch normalization to mitigate overfitting, and Activation Functions Sigmoid for gender classification and SoftMax for Age group classification. The dataset is divided into training and validation sets (70% and 30%, respectively), enabling the model to learn to map VGG-19 features to age and gender labels. To evaluate the performance of the model, metrics like accuracy, precision, recall, and F1-score are employed. The proposed model achieves an impressive 78.67% accuracy in predicting age and 97.02% accuracy in gender classification on the UTKFace dataset, outperforming traditional methods despite challenges posed by variations in lighting, pose, and expression. The robustness of our approach is evidenced by its capability to handle diverse gender representations.

Keywords: Age and Gender Classification; Deep Features; VGG-19; CNN; Identity Recognition; Human Facial Features.

Introduction:

Facial analysis has emerged as a pivotal area within computer vision, with utilization ranging from biometric systems to social media platforms. Understanding facial features enables machines to interpret human emotions, demographics, and behaviors, fostering advancements in various domains. The growing interest in facial analysis is fueled by the increasing availability of digital images and videos, along with the advancements in deep learning techniques that enable the extraction of rich features from facial data.

Within the broad domain of facial analysis, age and gender prediction from facial images have garnered remarkable attention due to their practical implications in diverse fields such as security, marketing, and personalized services [1]. Accurate prediction of age and gender from facial images facilitates targeted advertising, personalized recommendations, and enhanced user experiences in applications ranging from social media platforms to e-commerce websites. Additionally, age and gender prediction are crucial components in biometric systems and surveillance technologies, aiding in identity verification and monitoring [2].

Despite the progress made in age and gender prediction from facial images, challenges persist in achieving high accuracy and robustness, particularly in uncontrolled environments with variations in illumination, pose, and facial expression. Traditional methods often rely on handcrafted features and statistical models, which may not capture the complex patterns inherent in facial data. The most used handcrafted features are geometric, texture, and appearance-based. Geometric features analyze the shape and size of facial elements like the eyes, nose, and mouth to determine age and gender. Texture features use the patterns and wrinkles on the face to estimate age and classify gender [3]. Appearance-based features use the overall appearance of the face, including skin color, hair, and eye color, to estimate age and classify gender. Despite the significant progress made in these approaches, they still have some limitations. Geometric features are sensitive to facial expressions and head pose, which can affect the accuracy of age estimation and gender classification. Texture features are affected by lighting conditions, image resolution, and facial hair, which can also affect the accuracy of these tasks [4]. Appearance-based features are sensitive to the quality of the image and the presence of occlusions, which can reduce the accuracy of age estimation and gender classification. Additionally, variations in demographics and cultural differences pose challenges in developing generalized models for age and gender prediction.

This research aims to tackle the challenges of predicting age and gender from facial images by using deep learning techniques, specifically convolutional neural networks (CNNs), for both feature extraction and classification. Deep learning models have shown outstanding performance in various computer vision tasks, such as object recognition and image classification, by learning hierarchical representations from raw data automatically. By training CNNs on large-scale datasets containing diverse facial images, we aim to develop robust models capable of accurately predicting age and gender across different demographics and environmental conditions. Here our proposed deep model learns from unstructured data and has notably improved the accuracy and efficiency of age estimation and gender classification. The proposed model can handle variations in facial expressions, pose, and lighting conditions.

Novelty and Objectives:

The contributions of this research are twofold: Firstly, we present a comprehensive analysis of existing methods and challenges in age and gender prediction from facial images, highlighting the limitations of traditional approaches and the potential of deep learning techniques. Secondly, we propose novel methodology for age and gender prediction, leveraging advanced image processing techniques for data preprocessing and deep learning architecture for feature extraction and classification. Empirical evaluations on benchmark dataset were conducted to assess the effectiveness of our proposed methods. Results demonstrate strong performance across multiple metrics, including accuracy, robustness, precision, recall, and F1 score in both age and gender prediction tasks. Overall, this research contributes to advancing the state-of-the-art in facial analysis and recognition, with implications across various domains including security, marketing, and human-computer interaction.

Literature Review:

In this section, we explore the existing body of research and methodologies relevant to age and gender classification. Researchers have made significant strides in age and gender estimation from facial data, gait, and voice driven by advancements in DL and CV. Various techniques have been explored, ranging from traditional handcrafted features to state-of-the-art DNN. Notably, the availability of large-scale datasets has facilitated evaluation and benchmarking. A brief summary is provided in Table 1.

This study [5] present a DL-based solution This research focuses on predicting age and gender from facial images. The proposed method utilizes a CNN architecture comprising multiple convolutional and pooling layers, followed by fully connected layers. The classifier in their research is a Multi-Layer Perceptron (MLP). This research [6] demonstrates that substantial enhancements in task performance can be attained by leveraging learned representations through Convolutional Neural Networks (CNNs). In their investigation, they fine-tuned a pre-trained DCNN using the extensive UTKFace dataset and subsequently evaluated its effectiveness on the same dataset and achieved 90% accuracy for gender and MAE 6.2 for age. An artificially generated collection of facial images is presented [7] to balance the need for confidentiality with the requirement for a large number of facial images for deep learning training, the researchers utilized a Style-Generative Adversarial Networks (StyleGAN) generator to produce a set of facial images. They classify images based on the perceived gender by employing metrics that identify male and female facial characteristics. This manually annotated dataset which is being used to train three face gender classifiers, one customized network, and two previously trained networks (VGG16 and VGG19) based on visual geometric group designs. These three classifiers have been cross validated on two distinct datasets, including labeled photographs of real people. They use the UTKFace and Kaggle gender datasets for testing. Their experimental results indicate that training using AI-generated photographs produces equivalent performance with accuracies comparable to existing state-of-the-art systems that use actual images of persons. Each classifier's average classification accuracy is between 94% and 95%, comparable to other proposed approaches.

Sharma et al[8] in their proposed work, use improved CNN for AEGC. They evaluated the model on CACD, UTKFace, FG-NET, and IMDB-WIKI datasets. Their results for the UTKFace dataset showed accuracy for age and gender estimation of 94.01% and 99.86% respectively. This study[9] proposes a GRA-GAN technique that employs a Generative Adversarial Network (GAN) for image style transfer related to gender, race, and age. It integrates encoder and decoder characteristics through channel-wise and multiplication-based information fusion. Experimental results using four publicly available datasets (MORPH, AFAD, AAF, and UTK) show that this technique outperforms current SOTA methods. This research[10] focuses on categorizing gender and age from a range of face images—the proposed approach, which is built on a SegNet-based structure and an ML algorithm, provides good results. The redesigned Seg-Net structure and SVM enhanced overall accuracy for age and gender recognition. On multiple datasets including Adience, IOG, and FG-Net, their proposed solution beat existing technology, with an accuracy of 74.5%, 75.7%, and 92.48%, respectively, and Gender Classification with an accuracy of 88.3%, 95.1%, 94.1%, and 91.8%, respectively, when compared to Adience, IOG, FEI, and own datasets.

In this article [11] Three CNN network models are shown, each with a different design, comprising the number of filters and layers of convolution. These models were verified and tested using IMDB and WIKI datasets. According to the findings, CNN networks considerably improve system efficiency and recognition accuracy. In this study [12] the researchers suggest using Random Occlusion as a data augmentation method for age and gender recognition in facial images. They apply this technique to two different convolutional neural networks (CNNs) to classify individuals based on their age and gender: a modified version of Adience Net and VGG16. The research community frequently uses face recognition methods like (SIFT) and (SURF).

This study[13] develop a hybrid technique that combines Conventional Artificial Neural Networks (C-ANN) and Convolutional Neural Networks (CNN) for age and gender classification using decision fusion methods. A key aspect of this study is the integration of decisions from both neural networks to enhance the accuracy of age and gender predictions. To improve recognition accuracy, probabilistic decision fusion methods such as majority voting, Naive-Bayes combination, and sum rule were employed. The sum rule decision fusion technique outperformed existing methods by reducing the likelihood of misclassifying neighboring classes. A unique CNN model has been suggested for classifying unconstrained facial images based on age and gender in [14] Research. Several tests were carried out using the Adience dataset to assess the effectiveness of the suggested approach for identifying an individual to the proper age and gender. Experiment findings on the OIU-Adience dataset shows that their model outperforms previous studies on the same dataset in terms of classification accuracy, and demonstrating better performance. The proposed method achieves 84.8% accuracy in age group classification and 89.7% accuracy in gender classification.

Table 1: Summary of Existing State of the Art Techniques

Ref	Year	Dataset	Method	Acc %
[5]	2023	UTKFace	CNN+MLP	80.20 94.50
[6]	2023	UTKFace	DCNN	6.2 MAE 90.00
[7]	2023	AI-generated (Training). UTKFace (Testing)	StyleGAN VGG16+VGG19	94.50
[8]	2022	UTKFace,	Improved CNN	99.86 94.01
[9]	2022	UTKFace	GRA-GAN	6.12 MAE 90.82
[10]	2022	Adience	SegNet+SVM	74.50 88.30
[11]	2021	IMDB WIKI	CNN	86.20, 94.49 83.97, 93.56
[12]	2021	Adience	CNN VGG 16	89.00 92.40
[13]	2021	2000 Images from the Internet	GF+NN	86.10 98.40
[14]	2020	OIU-Adience	CNN(4C2FC)	84.80 89.70

Proposed Methodology:

Model Description:

The task of estimating the age group and gender of person based on facial image features poses significant challenges.

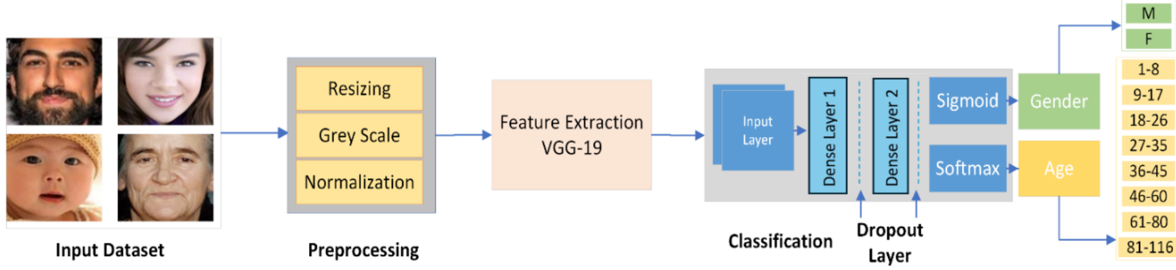

Figure 1, illustrates the flow diagram of the proposed method, detailing the entire process. It involves several key steps, including dataset collection, preprocessing of images, feature extraction from facial images using VGG19, and lastly age group and gender estimation applying a convolutional neural network (CNN). This framework aims to address the complexities of age and gender estimation using an integrated approach by leveraging the capabilities of both the VGG19 model and CNNs to achieve accurate and robust age and gender classification by using facial images.

Preprocessing:

The first step in the methodology is to preprocess the facial images. This involves resizing the images to a standardized dimension of 224x224 pixels, converting them to grayscale, and performing normalization to emphasize the quality of the images and reduce the impact of variations in lighting and color. We meticulously curated our dataset by selecting aligned and cropped face images from the UTKFace dataset, ensuring consistent positioning and eliminating background noise.

Feature Extraction using VGG-19:

The feature extraction process involves feeding preprocessed images into the VGG-19 network, renowned for its robust feature-detection capabilities across various levels of abstraction. The VGG-19 model architecture is characterized by a series of blocks, each block comprises of convolutional layers pursued by a max-pooling layer. The first block comprises

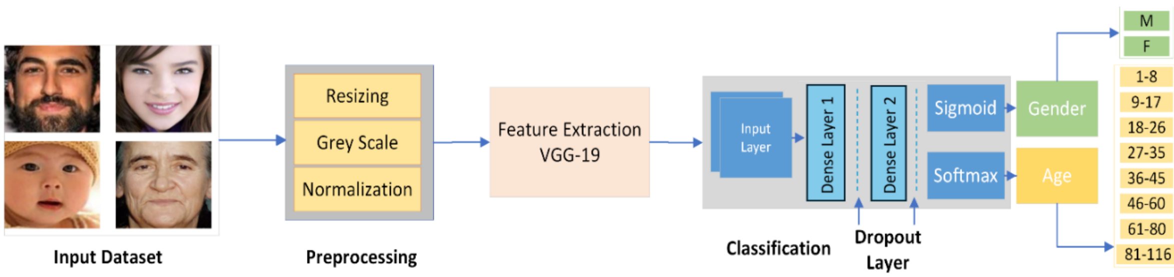

Figure 1: Proposed method Flow Diagram

two convolutional layers, each trailed by a rectified linear unit (ReLU) activation function, and then a max-pooling layer to downs ample the feature maps. This pattern is repeated for the second block. The third block in the VGG-19 architecture is slightly different, featuring four convolutional layers with ReLU activation functions followed by a max-pooling layer which is shown in Figure 2. This block is repeated three times, resulting a total of 16 convolutional layers and 5 max-pooling layers in the entire architecture. After the final max-pooling layer, the output is passed to a flattened layer, which reshapes the multi-dimensional feature maps into a one-dimensional vector. This vector serves as a compact representation of the extracted features from the input facial image, ready for further processing in subsequent layers of the neural network.

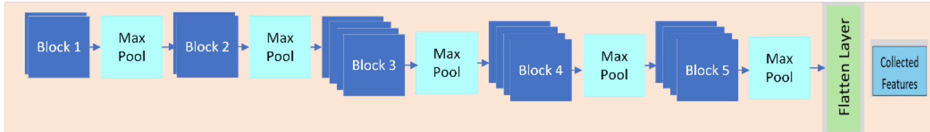

Figure 2: VGG-19 Architecture

Age and Gender Classification using CNN:

The classification phase involves employing a custom-designed CNN tailored for age and gender classification tasks. This CNN architecture is specifically configured to process one-dimensional feature vectors extracted by the VGG-19 model. CNN comprises several layers, including parallel branches of dense layers with ReLU activation functions, responsible for interpreting features and learning complex face patterns for classification. Dropout layers with a rate of 0.4 are implemented to prevent overfitting by randomly deactivating neurons during training. The network bifurcates into two branches for Gender and Age classification tasks.

For gender classification, A dense layer with a single neuron and a sigmoid activation function is employed, while for the age classification utilizes a dense layer with 8 neurons, each corresponding to a distinct age group. SoftMax activation function is applied to output a probability distribution across age groups. The entire classification pipeline is encapsulated by compiling the model with defined inputs and outputs.

We developed a custom CNN architecture for age and gender classification. The model processes one-dimensional feature vectors from the VGG-19 model. Key components include parallel branches for age and gender, dropout layers (with a rate of 0.4) to prevent overfitting, and specific activation functions (sigmoid for gender, SoftMax for age). Our hyperparameters were chosen through empirical experimentation. The bifurcation into separate branches enables the model to learn distinct features, while SoftMax activation aids accurate age predictions. Training occurs on deep features extracted from the UTKFace dataset, utilizing a batch size of 32 and training for 30 epochs and the detail of Training hyperparameters is provided in Table 2.

Table 2: Training Hyperparameters

Dense layers = 2	Units (256)	
Dropout Layers = 2	Dropout rate = 0.4	
Output Layers = 2	1 for Gender	Activation Function = Sigmoid
	2 for Age	Activation Function = SoftMax
Number of Epochs	30	
Batch Size	32	
Optimizer	adam	Learning Rate=0.001

Experiments and Results:

Dataset:

The UTKFace dataset is a rich and diverse repository comprising facial images spanning an extensive age spectrum, from infants to elderly individuals, with ages ranging from zero to 116 years. With over 23,000 meticulously annotated facial photographs, this dataset encompasses a broad array of age groups, genders, and ethnicities. Images are categorized as different age groups as toddlers and young children (1–8 years), pre-teens and teenagers (9–17 years), young adults (18–26 years), adults (27–35 years), mature adults (36–45 years), middle-aged individuals (46–60 years), seniors (61–80 years), and the elderly (81–116 years). Sample images taken from dataset are shown in **Error! Reference source not found.**. To effectively t rain the DL model, it is imperative to partition the dataset. This process involves randomly dividing the dataset images into two distinct subsets: a training set and a validation set. In this study, the dataset is split into 70% for training and 30% for validation. This partitioning strategy ensures that the model is exposed to a assorted range of images during the training process, thereby increasing its ability to generalize to unseen data.

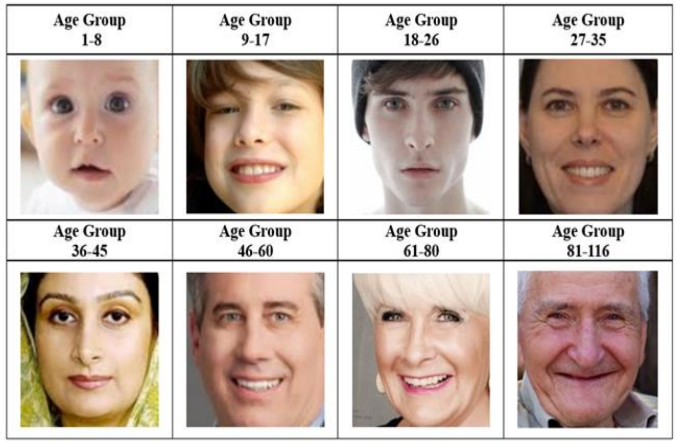

Figure 3: Sample images from UTKFace Dataset

Performance Metrics:

The evaluation of the proposed method is done using a variety of performance metrics to gauge the effectiveness of the models. These metrics include accuracy, loss, recall, specificity, F1-score, and area under the curve (AUC). Each metric provides valuable insights into the model's performance, aiding in its assessment and refinement.

Accuracy, defined as the ratio of correctly predicted instances to the total number of instances is mathematically described in equation (1), serves as a fundamental performance indicator. It offers a comprehensive assessment of the model's ability to accurately classify instances across all classes, reflecting its general predictive capability. Sensitivity, also referred to as recall, is described by equation (2) as the proportion of samples correctly identified as positive out of the total number of samples that are actually positive. This metric is also recognized as the true positive rate, reflecting the model's ability to accurately detect positive instances among all positive cases.

$$Accuracy\ Rate = \frac{True\ Positive + True\ Negative}{Total\ Positive + Total\ Negative} \quad (1)$$

$$Recall\ Rate = \frac{True\ Positive}{True\ Positive + False\ Negative} \quad (2)$$

The F1 score, represented in equation (3), is the harmonic mean of precision and recall. It provides a single metric to assess the balance between precision and recall, capturing both the model's ability to correctly identify positive instances (precision) and its ability to capture all positive instances (recall). This metric offers a thorough assessment of the model's performance, especially in situations with class imbalance or when both precision and recall hold equal significance. Equation (4) denotes precision, also known as positive predictive value, which signifies the proportion of samples correctly identified as positive out of the total number of samples predicted as positive. This metric measures the accuracy of positive predictions made by the model, offering insights into its ability to avoid false positives and make precise classifications.

$$F\text{-}1\,Score = 2 * \frac{(Precision * Recall)}{Precision + Recall} \quad (3)$$

$$Precision\ Rate = \frac{True\ Positive}{True\ Positive + False\ Positive} \quad (4)$$

Training and Validation Results:

The training and validation sets are partitioned with a ratio of 70% for training and 30% for validation, ensuring an adequate representation of images in both subsets. This random splitting of the dataset into two parts serves to maintain diversity in the images utilized for training and validation. By randomly selecting images for each set, we aim to capture a broad spectrum of variations present in the dataset, including different poses, expressions, lighting conditions, and other characteristics. This approach helps prevent bias and overfitting during model training, enabling more robust and generalizable performance on unseen data.

Figure 4 depicts the model accuracy for gender classification attained on proposed approach. It is obvious from the graph that the model's increasing ability to accurately predict gender as it processes more data during training. The maximum accuracy 97.02% achieved for gender classification. Figure 5 depicts the model's results in predicting different age groups accurately. The graph demonstrates a progressive increase in accuracy, suggesting that the model is becoming more adept at estimating ages correctly as it undergoes further training. For the age groups 78.67% accuracy was achieved.

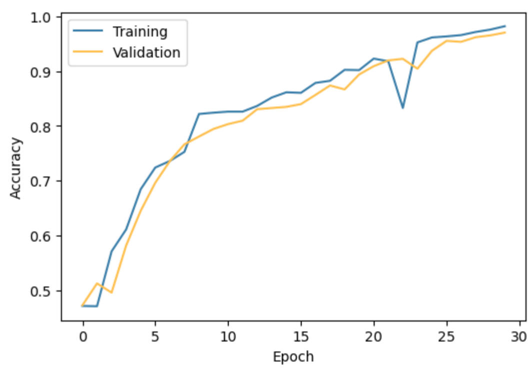

Figure 4: Gender Classification

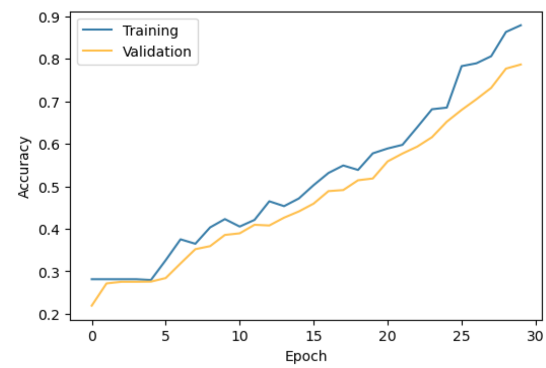

Figure 5: Age Group Result

In Figure 6 and Figure 7, we illustrate the comparative analysis of training and validation losses concerning gender and age classification. The validation loss metric serves as a gauge of the model's ability to generalize to unseen data, indicating how well it performs on data not used during training. On the other hand, the training loss metric quantifies the model's accuracy in replicating the outcomes observed in the training dataset. These loss metrics offer insights into the model's learning progress and its capacity to achieve optimal performance across both training and validation datasets.

Discussion:

To assess the effectiveness of our methodology, we compared the accuracy achieved with existing state-of-the-art approaches for age and gender classification. **Error! Reference source not found.** displays the outcomes of additional performance evaluation metrics for the proposed method. The results unequivocally illustrate that our proposed method yields significantly enhanced outcomes for age and gender classification using facial images compared to the results obtained by other existing methods. Figure 9 illustrate the Age group classification accuracy comparison while Figure 10 depicting Gender classification Accuracies comparison of our approach with existing methods. This comparative analysis emphasizes the superior performance and efficacy of our proposed approach in accurately predicting age and gender from facial images.

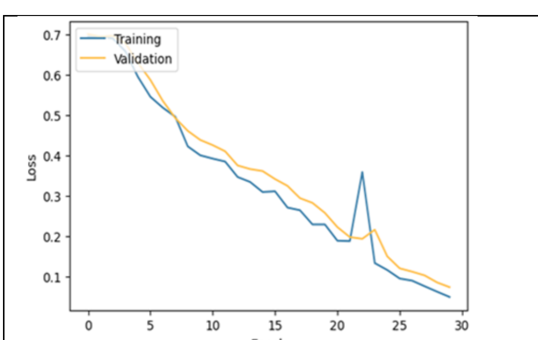

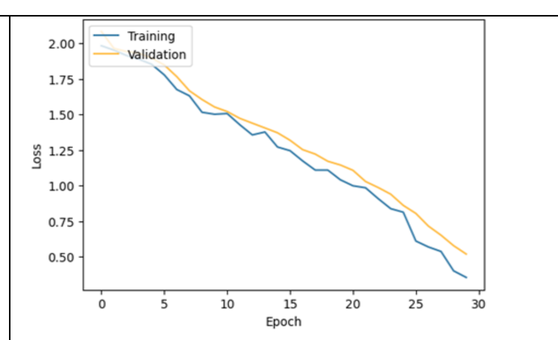

Figure 6: Loss for Gender	**Figure 7:** Loss for Age Groups

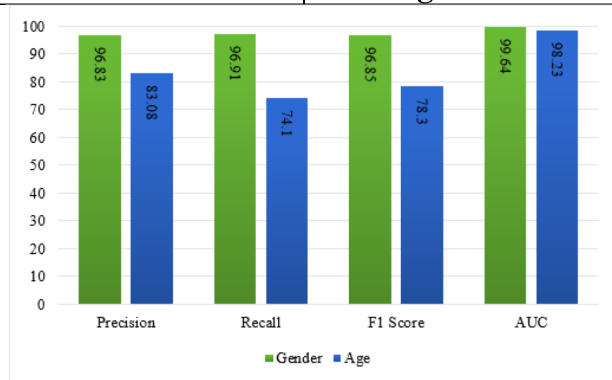

Figure 8: Performance Metrices for Age and Gender.

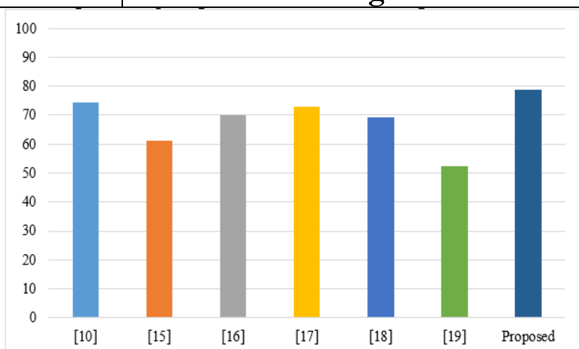

Figure 9: Comparison of Age Group prediction with SOTA Techniques.

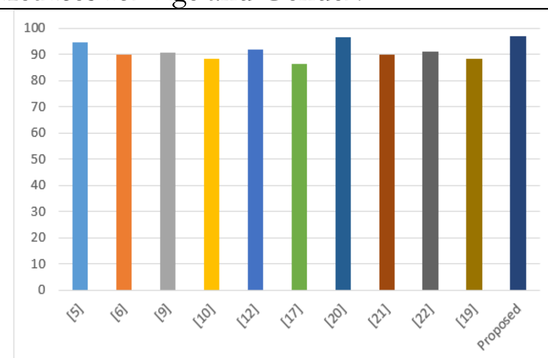

Figure 10: Comparison of Gender prediction with SOTA Techniques.

While our study demonstrates promising results, we acknowledge the possibility of further improvement. Currently, we utilize VGG-19 for feature extraction. However, considering alternative deep learning architectures for feature extraction may enhance the robustness and performance of our approach. We rigorously evaluated our model's performance and conducted statistical significance tests to validate improvements.

Conclusion:

This research introduces a robust methodology for age and gender prediction from facial images, employing deep features extracted using VGG-19 architecture and a custom CNN model for classification. The impressive accuracy rates achieved—78.67% for age prediction and 97.02% for gender classification on the UTKFace dataset—underscore the efficacy of our approach. By surpassing traditional techniques and demonstrating resilience against challenges like variations in lighting, pose, and expression, our method showcases the potential of deep learning models in facial analysis tasks. Furthermore, comparative analysis with existing methods highlights the superiority of our approach.

Future Work:

Future investigations may focus on extending this approach to larger and more diverse datasets, exploring additional fusion techniques for further enhancement, and addressing interpretability and fairness concerns to ensure ethical considerations in facial analysis applications.

Acknowledgement: We gratefully acknowledge the UTKFace dataset, generously provided on Kaggle, for its invaluable contribution to our research.

Author's Contribution. The study's conception and design involved contributions from all authors.

Conflict of interest. The authors declare that they have no conflicts of interest.

References:

[1] R. Angulu, J. R. Tapamo, A. O. J. E. J. o. I. Adewumi, and V. Processing, "Age estimation via face images: a survey," vol. 2018, no. 1, pp. 1-35, 2018.
[2] P. Terhörst, N. Damer, F. Kirchbuchner, and A. Kuijper, "Suppressing gender and age in face templates using incremental variable elimination," in 2019 International Conference on Biometrics (ICB), 2019, pp. 1-8: IEEE.
[3] A. Althnian, N. Aloboud, N. Alkharashi, F. Alduwaish, M. Alrshoud, and H. J. A. S. Kurdi, "Face gender recognition in the wild: an extensive performance comparison of deep-learned, hand-crafted, and fused features with deep and traditional models," vol. 11, no. 1, p. 89, 2020.
[4] X. Tan and B. J. I. t. o. i. p. Triggs, "Enhanced local texture feature sets for face recognition under difficult lighting conditions," vol. 19, no. 6, pp. 1635-1650, 2010.
[5] B. B. Mamatkulovich and H. A. J. E. S. H. Alijon o'g'li, "Facial Image-Based Gender and Age Estimation," vol. 18, pp. 47-50, 2023.
[6] T. P. Kancharlapalli and P. Dwivedi, "A Novel Approach for Age and Gender Detection using Deep Convolution Neural Network," in 2023 10th International Conference on Computing for Sustainable Global Development (INDIACom), 2023, pp. 873-878: IEEE.
[7] V. Raman, K. ELKarazle, P. J. C. S. S. Then, and Engineering, "Artificially Generated Facial Images for Gender Classification Using Deep Learning," vol. 44, no. 2, pp. 1341-1355, 2023.
[8] N. Sharma, R. Sharma, and N. J. W. P. C. Jindal, "Face-based age and gender estimation using improved convolutional neural network approach," vol. 124, no. 4, pp. 3035-3054, 2022.
[9] Y. H. Kim, S. H. Nam, S. B. Hong, and K. R. J. E. S. w. A. Park, "GRA-GAN: Generative adversarial network for image style transfer of Gender, Race, and age," vol. 198, p. 116792, 2022.
[10] S. Kumar, S. Singh, J. Kumar, K. J. M. T. Prasad, and Applications, "Age and gender classification using Seg-Net based architecture and machine learning," vol. 81, no. 29, pp. 42285-42308, 2022.
[11] M. K. Benkaddour, S. Lahlali, and M. Trabelsi, "Human age and gender classification using convolutional neural network," in 2020 2nd international workshop on human-centric smart environments for health and well-being (IHSH), 2021, pp. 215-220: IEEE.
[12] C.-Y. Hsu, L.-E. Lin, C. H. J. M. T. Lin, and Applications, "Age and gender recognition with random occluded data augmentation on facial images," vol. 80, pp. 11631-11653, 2021.
[13] J. Rwigema, J. Mfitumukiza, K. J. B. S. P. Tae-Yong, and Control, "A hybrid approach of neural networks for age and gender classification through decision fusion," vol. 66, p. 102459, 2021.
[14] O. Agbo-Ajala and S. Viriri, "Face-based age and gender classification using deep learning model," in Image and Video Technology: PSIVT 2019 International Workshops, Sydney, NSW, Australia, November 18–22, 2019, Revised Selected Papers 9, 2020, pp. 125-137: Springer.
[15] Hosseini, S., et al. Age and gender classification using wide convolutional neural network and Gabor filter. in 2018 International Workshop on Advanced Image Technology (IWAIT). 2018. IEEE.
[16] Das, A., A. Dantcheva, and F. Bremond. Mitigating bias in gender, age and ethnicity classification: a multi-task convolution neural network approach. in Proceedings of the european conference on computer vision (eccv) workshops. 2018.

[17] Hassan, K.R. and I.H. Ali. Age and Gender Classification using Multiple Convolutional Neural Network. in IOP Conference Series: Materials Science and Engineering. 2020. IOP Publishing.
[18] Khan, K., et al., A multi-task framework for facial attributes classification through end-to-end face parsing and deep convolutional neural networks. 2020. 20(2): p. 328.
[19] Duan, M., et al., A hybrid deep learning CNN–ELM for age and gender classification. 2018. 275: p. 448-461.
[20] Singh, A., V.K.J.N.C. Singh, and Applications, A hybrid transformer–sequencer approach for age and gender classification from in-wild facial images. 2024. 36(3): p. 1149-1165.
[21] Mustapha, M.F., et al. Improving the Accuracy of Gender Classification Based on Skin Tone Using Convolutional Neural Network: Transfer Learning (CNN-TL). in International Conference on Robotics, Vision, Signal Processing and Power Applications. 2024. Springer.
[22] Uysal, M. and M.F.J.D.Ü.M.F.M.D. Demiral, Gender and Age Estimation By Image Processing. 15(1): p. 49-59.