Grenze International Journal of Engineering and Technology, June Issue

Dissecting Facial Features: A Deep Learning Synthesis for Age and Gender Prediction

Dr.Ch. Rajendra Babu[1], Thalam Jaya Dharani[2], Battula Sai Durga Rao[3] and Shaik Bibi Ayesha[4]
[1-4]Lakireddy Bali Reddy College Of Engineering, Mylavaram, India
Email: chikkalarajendra@lbrce.ac.in, thalamdharanibindu18@gmail.com, Saidurgarao409@gmail.com, shaikbibiayeshabibiayesha8179@gmail.com

***Abstract*—Age and gender detection is crucial for enhancing user experience and security in applications like personalized content delivery and demographic analysis. However, current methods often struggle with accuracy due to diverse facial features and challenging imaging conditions. This research proposes a deep learning-based approach that uses ResNet-50 and Convolutional Neural Networks (CNN) to address these issues. We trained our model on the UTKFace dataset, leveraging residual networks to enhance feature extraction, resulting in higher accuracy and efficiency in real-world scenarios. According to experimental results, ResNet-50 performs better than conventional CNN models in terms of accuracy and processing speed, which qualifies it for real-time applications. Our approach provides a scalable, reliable solution to age and gender detection, with the potential for broad deployment in various industries.**

***Index Terms*— Age Detection, Gender Classification, Deep Learning, CNN, ResNet-50, UTKFace Dataset, Facial Recognition, Demographic Analysis, Real-Time Applications, Image Processing.**

I. INTRODUCTION

Age and gender detection has emerged as a critical component in various applications across multiple industries, including personalized content delivery, targeted advertising, demographic analysis, and security systems. The ability to automatically recognize individual characteristics through facial analysis enables organizations to offer tailored experiences, improving user engagement and operational efficiency. This technology plays a vital role in providing insights into user demographics, which businesses can leverage to design products and services better suited to their audiences' needs. Additionally, in sectors like security, accurate age and gender detection aids in surveillance, identity verification, and access control. The absence of automated age and gender detection systems poses several challenges. Without such systems, businesses often rely on manual or generic methods to segment their audiences, leading to inefficiencies and reduced personalization. For instance, advertising campaigns may fail to reach the intended demographic, or content recommendations may not align with user preferences, resulting in diminished user satisfaction. In security, the inability to detect individual characteristics could lead to vulnerabilities in monitoring and authentication processes. These challenges underscore the importance of developing robust and accurate age and gender detection systems that can operate effectively across diverse scenarios. The applications of age and gender detection span a wide range of domains. In marketing, these systems enable targeted campaigns by identifying demographic segments accurately. In

Grenze ID: 01.GIJET.11.2.557

entertainment, they enhance user experiences by recommending content aligned with viewers' age groups and preferences. In healthcare, they assist in patient profiling and treatment planning. In security, they contribute to advanced surveillance systems and fraud prevention. The integration of age and gender detection technology thus offers significant benefits, making it a cornerstone for innovation across industries. Despite the potential, accurate age and gender detection faces numerous challenges. Human facial diversity, encompassing variations in ethnicity, skin tone, facial expressions, lighting conditions, and image quality, complicates prediction accuracy. Traditional machine learning models that depend on handcrafted features struggle to generalize across diverse datasets, often leading to issues with precision and reliability in real-world applications. These limitations call for advanced methodologies that can handle the complexity of facial attributes effectively. Recent advancements in deep learning, particularly Convolutional Neural Networks (CNNs), have altered facial analysis.CNNs are incredibly effective for applications like age and gender recognition because, in contrast to conventional methods, they automatically learn hierarchical features from raw picture data. CNNs perform better than conventional techniques, but they have trouble processing big datasets or obtaining the fine-grained characteristics needed for accurate classification. To address these challenges, this Study employs ResNet-50, a state-of-the-art deep learning architecture built upon CNNs, incorporating residual learning to overcome common issues like vanishing gradients in deep networks. ResNet- 50’s design enables it to focus on learning complex and informative patterns, making it well-suited for tasks requiring detailed feature extraction, such as age and gender detection. The UTKFace dataset, a sizable and varied collection of annotated photos with features related to age, gender, and ethnicity, is used in this study to train and assess the model. This compares and contrasts the CNN and ResNet-50 architectures, highlighting their advantages and disadvantages. Through the use and analysis of these methods, the research seeks to improve the precision and effectiveness of age and gender categorization, offering a real-time, scalable solution appropriate for a range of sectors, such as marketing, entertainment, security, and healthcare. This work not only demonstrates the application of cutting-edge deep learning techniques but also highlights their potential to overcome traditional challenges, providing a foundation for future advancements in facial intelligence systems.

Fig.1. Age and Gender detection

II. RELATED WORK

This study explores age and gender detection using deep learning, particularly CNN and hybrid models, for real-world applications. Hybrid CNN-RNN Architecture for Gender and Age Detection (2024) integrates CNN with RNN to leverage spatial and sequential data, achieving high accuracy in healthcare and security. Gender Prediction and Age Estimation using CNN (2021) *and* Real-Time Age and Gender Detection Using Deep Learning Techniques (2020) enhance reliability in customer service and enterprise settings, employing SSD and MobileNetV2 for accurate detection in high-traffic areas like airports. Age and Gender Classification using CNN (2022) improves accuracy by linking image-derived demographic data with form inputs for better personalization in digital marketing. *GRA_Net (2023)* incorporates selective attention mechanisms into Residual Attention Networks for applications in retail and targeted advertising. Gender and Age Estimation Based on Deep Learning Techniques (2023) reviews demographic diversity challenges and solutions for masked faces, relevant to biometric security. Automated Detection of Gender from Face Images (2020) enables real-time classification of age, gender, and emotion in video streams. Real-Time Age and Gender Recognition using CNN (2023) achieves 84% accuracy in age estimation and 88.67% in gender classification, benefiting retail and entertainment. Age and Gender Prediction using OpenCV (2022*)* examines 3D face detection for automatic attendance and access control in educational settings. *Real-Time Gender Classification using CNN (2021)* focuses on optimizing CNN models for identity verification and targeted content delivery. Age Group and Gender Estimation in the Wild with Deep RoR Architecture (2021) introduces Residual Networks of Residual Networks (RoR) for processing unstructured image data, excelling in surveillance and smart city applications. These studies collectively highlight CNN-based and hybrid architectures for age and gender detection, tackling real-world challenges while advancing real-time processing, cross-verification, and multi-functional classification.

III. METHODOLOGY AND SOLUTION APPROACH

A. Dataset Information

The proposed system utilizes the UTKFace dataset, a comprehensive collection of over 20,000 facial images labeled with age, and gender. Each image represents an individual's face, spanning a diverse age range from 0 to 116 years, which ensures robust age prediction across various demographic groups. The dataset includes metadata for each image, specifying age and gender labels, making it ideal for training a model aimed at age and gender detection. However, for this study, the focus is on the age and gender attributes alone. Before training, the dataset undergoes essential preprocessing to ensure it's ready for model training.

Fig.2. Data Set For The Age and Gender Detection

B. Data Preprocessing

Preprocessing the data is an essential step that improves the efficiency and accuracy of the model. To preserve consistency and make it easier to integrate the UTKFace dataset photos with deep learning models such as CNN and ResNet-50, the images are downsized to a standard 224x224 pixel resolution. In order to accelerate model convergence during training, images are also normalized to values between 0 and 1. Data augmentation is used to increase the generalization and robustness of the model. To replicate a variety of real-world situations, this incorporates transformations like zooming, random rotations, and horizontal flipping. In order to provide efficient model learning for multi-class classification tasks, age values are classified into distinct age categories (child, teenager, adult, and senior) for label handling, and gender labels are encoded as binary (male and female).

C. Model Architecture

Convolutional Neural Networks (CNN) and ResNet-50 are two deep learning models that the system uses to identify age and gender from facial photos. It combines rapid feature extraction with sophisticated deep learning capabilities for comprehensive feature recognition, utilizing the strengths of both models to produce predictions that are accurate and efficient. Convolutional layers are used by the CNN model to process facial images and extract important properties including forms, edges, and textures. In order to identify fundamental visual patterns, such as facial contours, skin textures, and structural elements like eyes and nose, filters are applied to the input image. The spatial dimensions of the feature maps are decreased by using max-pooling layers, which lowers computational complexity without sacrificing significant features. Following dimensionality reduction and feature extraction, the features are passed through fully connected layers to make the final predictions for age and gender. Age prediction is handled as a regression task, where the model predicts an age value or assigns the image to a specific age range. Gender prediction is a binary classification task, where the model classifies the image as either male or female. On the other hand, ResNet-50 uses residual blocks, which add skip connections that let the network avoid specific levels, to overcome the difficulties of very deep networks. This ensures that the network can learn more effectively even in very deep layers, addressing the vanishing gradient problem. These residual connections help the model capture complex facial features and maintain learning efficiency as the model depth increases. The 50 layers in ResNet-50 allow the model to learn intricate, high-level features such as wrinkles, subtle skin changes, and other facial attributes that are crucial for distinguishing between age groups and genders. The bottleneck architecture in ResNet-50 optimizes computation, ensuring the model learns from complex features without being computationally expensive. Instead of fully connected layers at the end of convolutional layers, ResNet-50 uses global average pooling, which reduces feature map dimensions and produces a compact representation of the input image, which is then passed to the final layer for predictions. Together, the CNN model efficiently extracts the most important features from the input images, while ResNet-50 enhances these features by learning complex, deeper patterns crucial for accurate age and gender classification. This hybrid approach allows the system to be efficient and accurate, capable of delivering high-quality predictions even with large and complex datasets. This integration of CNN and ResNet-50 provides a robust solution for age and gender detection. CNN serves as a feature extractor, ensuring quick processing, while ResNet-50 enhances the extracted features for deeper analysis, providing more accurate predictions. The system is designed to work in real time, making it suitable for applications like face-based authentication systems or social media platforms. Using both models ensures a balance of speed and accuracy, with the ability to detect subtle facial features that are important for age and gender classification.

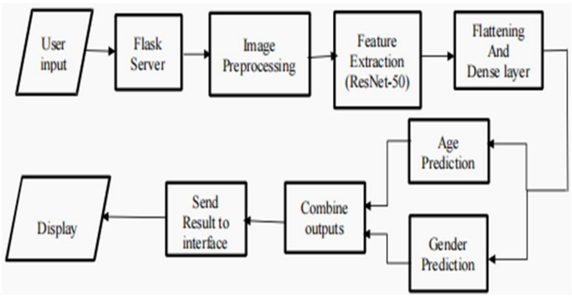

Fig.3.Proposed Model Architecture

D. Training and Evaluation

The data-set is divided into sets for testing, validation, and training. The models are trained utilizing the Adam optimizer for faster convergence and the categorical cross-entropy loss function to maximize classification accuracy. Early stopping tracks validation accuracy and stops training when performance reaches a plateau in order to avoid over fitting. To ensure a thorough evaluation of model capabilities, performance is assessed using accuracy, precision, recall, and F1-score. To guarantee that predictions can be made quickly in real-time scenarios—a crucial component for user satisfaction in real-world applications—inference time is also monitored.

E. System Deployment

After the models have been successfully trained, they are integrated into a Flask-based web application. Flask, a lightweight Python framework, is chosen because it provides an efficient and simple way to deploy machine learning models in a web environment. The integration allows users to easily upload images through the web interface and receive instant feedback on predicted age groups and gender. Each image is processed through a well-defined model pipeline, leveraging the trained machine learning models to generate near-instantaneous predictions, which are displayed immediately on the user interface. This ensures a smooth and seamless user experience, making the application highly responsive. The real-time capability of the application makes it highly valuable for various use cases, such as personalized content recommendations, targeted advertising, and security applications. For instance, the system
can enhance content recommendations by accurately identifying the user's age group and gender, allowing for a more tailored and engaging experience. Similarly, advertisers can leverage these predictions to serve highly relevant ads to different demographic groups, increasing the effectiveness of their campaigns. In security contexts, the application’s ability to predict age and gender can contribute to identity verification and surveillance systems. Additionally, the online application has a straightforward, easy-to-use interface. Users may input photographs rapidly, and the program responds immediately by processing them and making accurate predictions in seconds. This makes the system perfect for real-time applications where speed and usability are

essential.

IV. RESULTS

ResNet-50 outperforms traditional CNNs in gender classification (94.82% accuracy) and age prediction (MAE 2.7023). Validation results show 87.29% accuracy and MAE of 3.8581, indicating strong generalization. Residual connections enhance feature extraction, improving accuracy across diverse facial images. Its fast processing makes it ideal for real-time applications like surveillance and personalization.

Fig.4. Results Of Model

A. Deployment

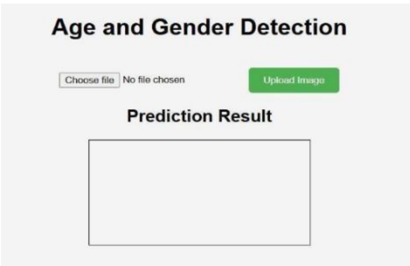

Fig.5. Interface of the Deployment

Fig.6.Sample Result of the Prediction

Table 1 compares the performance of different models, including ResNet-50, CNN, DLDL-v2, and TinyAgeNet, for age and gender prediction using the UTKFace dataset. ResNet-50 stands out with the highest accuracy of 94.82%, along with strong precision, recall, and F1-score for both male and female predictions, indicating its effectiveness in the task. The CNN model performs moderately with an accuracy of around 80%, showing good results but not as robust as ResNet-50. Both TinyAgeNet and DLDL-v2 show good accuracy (90% and unspecified, respectively), though their performance details are not fully provided. Overall, ResNet-50 performs best in terms of accuracy and other evaluation metrics, making it the most reliable choice for this task**.**

TABLE I: TABLE OF THE MODEL ANALYSIS

Model	**Accuracy**	**Precision**	**Recall**	**F1-Score**	**Dataset**
ResNet-50	**94.82%**	**~89.2%**	**~78.7%**	**~83.6%**	**UTKFace**
CNN	**~80%**	**~80%**	**~70%**	**~77%**	**UTKFace**
DLDL-v2 / TinyAgeNet	**90%**	**Not specified**	**Not specified**	**Not specified**	**UTKFace**

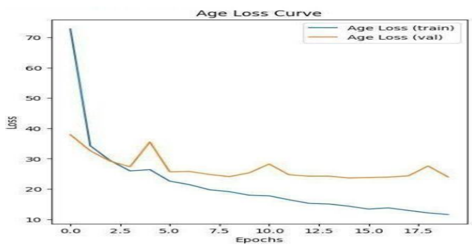

Fig.7. Age Loss Curve

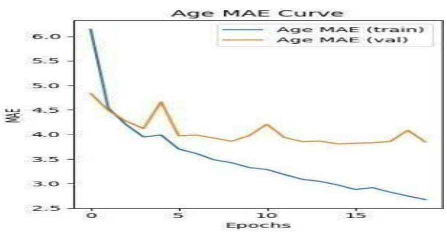

Fig.8.Age MAE Curve

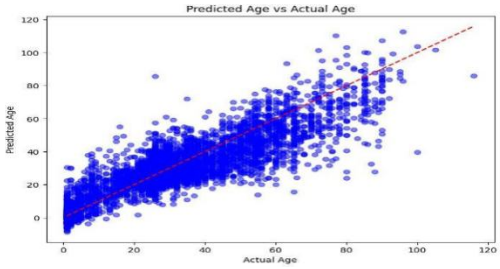

Fig.9. Prediction Age Vs Actual Age

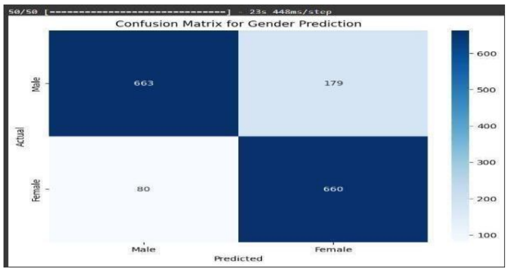

Fig .10. Confusion Matrix for Gender Prediction

V. CONCLUSIONS

In Summary, this study highlights the impressive capabilities of the ResNet- 50 models in age and gender detection, surpassing traditional CNNs with improved accuracy and speed, which is essential for real-world use cases like personalized recommendations and demographic analysis. Moving forward, broadening the dataset, adding features like emotion detection, optimizing for edge devices, and enhancing real-time performance could make the model even more robust and adaptable. Tailoring it to specific industry needs with transfer learning will also expand its potential applications, making it a flexible and powerful solution across various fields.

REFERENCES

[1] Arain, A., Ray, B., Singh, P.K., Ahmadian, A., Senu, N., & Sarkar, R. (2021). GRA_Net: A Deep Learning Model for Classification of Age and Gender From Facial Images. IEEE Access, 9, 85672-85689. DOI: 10.1109/ACCESS.2021.3087665.

[2] Zaghbani, S., Boujneh, N., &Bouhlel, M. (2018). Age Estimation Using Deep Learning. Computers & Electrical Engineering, 68, 337-347. DOI: 10.1016/j.compeleceng.2018.04.012.

[3] Gao, B.B., Zhou, H.Y., Wu, J., & Geng, X. (2018). Age Estimation Using Expectation of Label Distribution Learning. DOI: 10.24963/ijcai.2018/99.

[4] Li, S., & Cheng, K.T. (2019). Facial Age Estimation by Deep Residual Decision Making.

[5] ElKarazle, K., Raman, V., & Then, P. (2022). Facial Age Estimation Using Machine Learning TechniquesAn Overview. Big Data and Cognitive Computing, 6(4), 128. DOI: 10.3390/bdcc604012.

[6] Chandaliya, P., & Nain, N. (2022). AW-GAN: Face Aging and Rejuvenation Using Attention with Wavelet GAN. Neural Computing and Applications, 35, 1-15. DOI: 10.1007/s00521-022-07721-4.

[7] Suhasini, M., Apoorva, T.C., Vinay, M., Sunil, M.S., & Harshitha, M.N. (n.d.). Age and Gender Classification Using Convolutional Neural Network. Maharaja Institute of Technology, Thandavapura.

[8] Santoshi, K.N., Peddinti, S., Korada, K.S., Patchigoolla, A., &Hanumanthu, K.V. (2022). Gender and Age Detection with OpenCV. Department of Computer Science Engineering, GITAM Deemed to be University, Visakhapatnam, India-530045.

[9] Mishra, P., Singh, N., & Chavan, P. (2021). Real-Time Gender Classification Using Convolutional Neural Network. ITMWeb of Conferences, 40, 03047. DOI: 10.1051/itmconf/20214003047.

[10] Karnati, M., Mahidhar, D., Srija, D., Kalluru, K., & Sudhakar, A. (2024). Age and Gender Prediction. IRACSTInternational Journal of Computer Networks and Wireless Communications (IJCNWC), 14(2), 1-6.

[11] Awad, A., Sontakke, S., Narole, N., Nikalje, V., &Khaple, A. (2023). Real-Time Age and Gender Recognition Using CNN.International Journal of Research Publication and Reviews, 15(1).

[12] Salihbasi, A., &Orehovacki, T. (2019). Development of Android Application for Gender, Age, and Face RecognitionUsing OpenCV. MIPRO 2019, 20 May. DOI: 10.1109/MIPRO.2019.8757020.

[13] Xie, J.C., & Pun, C.M. (2019). Age and Gender Classification Using Convolutional Neural Networks. IEEE Transactions on Information Forensics and Security. DOI: 10.1109/TIFS.2019.2957003.