Novel Gender And Age- Based Detection Technique for Facial Recognition System

Pratham Gupta
School of Computer Science and Engineering, Vellore Institute of Technology, Chennai, Tamil Nadu, India

Amutha S
School of Computer Science and Engineering, Vellore Institute of Technology, Chennai, Tamil Nadu, India

Dhanush R
School of Electronics and Engineering, Vellore Institute of Technology Chennai, India

Heshalini Rajagopal
Department of Electrical & Electronic Engineering, Mila University, No. 1, Persiaran MIU, 71800 Putra Nilai, Negeri Sembilan Darul Khusus, Malaysia

Email: mail.guptapratham@gmail.com, dhanush.r@vit.ac.in, amutha.s@vit.ac.in, heshalini.rajagopal@mila.edu.my

Abstract

This paper introduces gender and age-based classification approaches in facial recognition systems, addressing challenges posed by diverse demographic characteristics. The model learns typical facial features and identifies deviations from them by employing unsupervised detection methods using autoencoders, enhancing robustness and generalization across populations. Ethical considerations are discussed, emphasizing the importance of fairness and bias mitigation in facial recognition. Experimental results demonstrate the method's effectiveness in handling biases in traditional supervised approaches. This research contributes a novel technique while highlighting the ethical implications of facial recognition. Gender and age-based detection methods improve system reliability in real-world scenarios with diverse demographics. These findings are relevant to researchers, developers, and policymakers navigating the intersection between facial recognition and ethical AI, promoting more responsible and inclusive technology.

Keywords: Facial recognition, Supervised learning, Age and gender classification, Neural network, Demographic characteristics, Generalization, Facial images

1. Introduction

Gender and age-based detection methods for facial images have garnered significant attention in computer vision research due to their wide spectrum of applications ranging between security and healthcare. In the real world, the precise identification of facial images based on gender and age presents unique challenges, including variations in lighting, facial expression, and pose. Such challenges were addressed over the years using diverse methodologies, with a notable shift towards deep learning concepts such as convolutional neural networks (CNNs). This introduction sets the stage for understanding the evolution of gender and age detection algorithms, highlighting key studies and advancements in the field.
Recent advancements in deep learning have revolutionized the field of demographic classification, particularly in age and gender estimation from facial images. Levi and Hassner's groundbreaking work in 2015 [1] marked a pivotal moment by demonstrating the efficacy of deep Convolutional Neural Networks (CNNs) in achieving high accuracy in demographic predictions. Their innovative use of CNNs showcased the potential of hierarchical feature extraction for capturing subtle patterns in facial data, setting a foundational benchmark that inspired subsequent research in the field. Building on Levi and Hassner's foundational work [1], Xing et al. [2] expanded the scope of demographic classification by integrating race classification into their model through multi-task learning. This approach allowed the model to handle multiple demographic attributes simultaneously, improving its robustness and versatility. Multi-task learning leverages shared representations, which enhances the model's capability to generalize across different tasks. Eidinger, Enbar, and Hassner [3] tackled the challenges in age and gender classification by using dropout-SVMs and robust face alignment techniques. Dropout-SVMs helped in improving model generalization by preventing overfitting, while robust face alignment techniques ensured accurate localization of facial features under various conditions. Kelliher's [4] comparative study highlighted the superiority of CNNs over traditional SVMs for complex visual tasks such as age and gender classification, systematically demonstrating how deep learning models outperform traditional machine learning approaches in handling the intricacies of visual data. Lee et al. [5] focused on

optimizing feature extraction techniques in gender classification, introducing novel approaches such as improved pre-processing steps, advanced network architectures, and refined training methods, which significantly enhanced prediction accuracy. Liu et al. [6] proposed a CNN-based model specifically designed for automatic age classification, showcasing robust performance across various age groups and demographic variations, highlighting the adaptability and scalability of deep learning techniques. Sundararajan and Woodard's [7] research enhances biometric systems by integrating deep learning, specifically improving age and gender classification accuracy using convolutional neural networks (CNNs). Their approach tackles challenges like facial variability, crucial for reliable authentication and identification in practical applications. CNNs excel in learning facial features despite lighting, expressions, and aging, providing robustness. This advancement promises more secure access control, reliable identity verification in finance, and effective law enforcement applications. Overall, their study signifies a significant stride in leveraging CNNs to bolster biometric system reliability and security, addressing critical limitations and enhancing real-world deployment efficacy. Hamdi and Moussaoui's [8] comparative analysis rigorously examined the efficacy of deep learning models versus traditional methods in age, gender, and ethnicity identification. Their study meticulously evaluated various approaches, highlighting deep learning's superiority in accuracy and robustness. By contrasting these models, they elucidated nuanced strengths and limitations, offering valuable insights for biometric system developers and researchers. The findings underscore deep learning's capability to handle complex data patterns and variability inherent in demographic attributes like age, gender, and ethnicity. Sharma et al. [9] showed better performance in age and gender classification through a pioneering approach that integrated advanced training techniques and utilized extensive datasets. By harnessing these strategies, Sharma et al. significantly enhanced classification accuracy and robustness, surpassing previous benchmarks. Their achievement underscores the importance of methodological innovation and dataset scale in advancing biometric systems, offering insights into effective strategies for improving age and gender classification accuracy in practical applications. Wang et al. [10] developed a sophisticated age estimation model that focused on capturing intricate ageing patterns, particularly in scenarios with limited visual cues. Their model aimed to enhance precision in demographic analysis by addressing challenges such as subtle facial changes over time. By leveraging advanced algorithms and robust data analysis techniques, Wang et al. achieved significant improvements in age estimation accuracy, contributing to more reliable and nuanced demographic assessments in various practical contexts. Meinedo and Trancoso [11] emphasized age and gender classification's pivotal role in forensic applications, particularly in contexts requiring precise demographic analysis like identifying child abuse material. Their focus underscores the importance of accurate biometric techniques in sensitive legal and investigative settings. By advancing methods for age and gender classification, Meinedo and Trancoso contribute to enhancing forensic tools' efficacy, aiding law enforcement in addressing critical issues related to child protection and criminal investigations. Saxena, Singh, and Singh [12] stressed the importance of robust demographic detection for enhancing security and social services, underscoring how accurate demographic predictions play a crucial role in improving safety and service delivery across various domains. Shaker and Al-Khalidi [13] demonstrated robust gender and age detection suitable for surveillance and access control systems, emphasizing the practical utility of accurate demographic classification in enhancing security and operational efficiency. Nada et al. [14] developed a solution for reliable demographic verification from profile photos using advanced image processing techniques, crucial for maintaining accuracy in online platforms. Ghosh and Bandyopadhyay [15] proposed an integrated method for gender classification and age detection, emphasizing feature extraction, selection, and classification techniques. Their approach aimed to enhance accuracy and efficiency across varied datasets, offering a robust framework for biometric applications requiring reliable demographic analysis. Brandao [16] focused on mitigating age bias in pedestrian detection algorithms to enhance fairness and inclusivity in demographic prediction models. Their work aimed to reduce disparities in algorithmic accuracy across different age groups, contributing to more equitable outcomes in automated systems relying on demographic analysis. Kumbhar and Shingare [17] emphasized the importance of comparative studies in deep learning algorithms for age and gender classification, providing a detailed evaluation of various models to guide future improvements. Wang, Ali, and Angelov [18] introduced an approach for gender and age classification crucial for detecting anomalous human behavior, contributing to the development of systems capable of identifying and responding to unusual activities. Karahan et al. [19] engineered machine learning algorithms specifically to elevate real-time accuracy in age and gender classification tasks. Their research underscores the significance of achieving robust performance under real-world conditions, enhancing the reliability and applicability of biometric systems in practical settings such as security and surveillance. Zaman et al. [20] expanded into multimodal biometric systems, enhancing voice-based demographic predictions by leveraging deep learning's versatility to improve accuracy across multiple modalities, highlighting the benefits of integrating multimodal data for comprehensive demographic predictions.

This paper introduces a gender- and age-based classification method in facial recognition systems. Traditional approaches often rely on supervised learning, which may encounter challenges in handling diverse demographic characteristics. Our proposed technique employs unsupervised detection to identify deviations from the norm, providing a more robust solution. We utilize autoencoders, a type of neural network well-suited to learn the inherent features of facial images. This research not only contributes a novel detection technique to the field but also emphasizes the importance of ethical AI in facial recognition. The application the novel detection techniques using gender and age classification enhances the reliability of facial recognition systems, making them more suitable for real-world scenarios where diverse demographic characteristics are prevalent.

2. Methodology

2.1. Dataset Description

Zhang .et. al [21] [22] aggregated large-scale facial image datasets with over 20,000 face images with annotations of age, gender, and ethnicity.

Here's a breakdown of the key components:

i. Images: The dataset contains over 20,000 RGB images of faces. These images vary in resolution and quality but generally provide a diverse set of facial characteristics.

ii. Age Annotation: Each face image is annotated with the subject's age. This annotation spans from 0 to 116 years old, making it one of the few datasets to cover such a broad age range.

iii. Gender Annotation: Alongside age, each image is also annotated with the subject's gender, indicating whether the face belongs to a male or female.

iv. Ethnicity Annotation: The dataset includes annotations for the ethnicity of each subject. This provides additional information for researchers interested in studying facial recognition across different ethnic groups.

v. Variability: The dataset captures various factors influencing facial appearance, including pose, expression, illumination, ethnicity, and age. This variability makes it suitable for training and evaluating algorithms robust to real-world conditions.

vi. Research Applications: UTKFace dataset is commonly used for research in age estimation, gender classification, and ethnicity recognition. Additionally, it can be utilized for tasks like facial recognition, facial expression analysis, and age progression/regression.

vii. The UTKFace dataset has been widely used in the computer vision and machine learning communities due to its large size, diversity, and comprehensive annotations, making it a valuable resource for researchers working on facial analysis tasks.

2.2. Data Preprocessing

Data augmentation techniques involve artificially increasing the diversity of the training dataset by applying transformations such as rotation, flipping, or scaling to the input images.

2.3. CNN for Multi Output Classification Model (CNN_MOCM)

The CNN_MOCM that we have made, constructs a convolutional neural network (CNN) with multiple convolutional and fully connected layers for gender and age prediction, a proposed method could involve the following enhancements:

· Multi-Task Learning (MTL): MTL involves jointly training the model for multiple related tasks, in this case, gender and age prediction. By sharing information between tasks during training, MTL can lead to improved generalization and performance compared to training separate models for each task.

- Attention Mechanisms: Attention mechanisms enable the model to dynamically focus on relevant regions or features of the input data. In the context of facial recognition, attention mechanisms can help the model prioritize important facial features for gender and age prediction, such as eyes, nose, or mouth regions.
- Transfer Learning: Transfer learning involves leveraging knowledge learned from pre-training on a large dataset (e.g., ImageNet) and fine-tuning the model on the specific gender and age prediction task. By initializing the model with pre-trained weights, transfer learning can accelerate the training process and boost performance, especially when the target task has limited training data.
- Regularization: The issue of overfitting could be avoided by adding a penalty term to the loss function. Thereby bypassing the use of large weights in the model. By penalizing complex models, regularization encourages simpler and more generalizable representations, leading to improved performance on unseen data.
- Hyperparameter Tuning: Hyperparameter tuning involves systematically searching for the optimal values of hyperparameters (e.g., learning rate, dropout rate) that govern the training process. By fine-tuning these hyperparameters, we can optimize the model's performance and convergence speed, leading to better results on the gender and age prediction task.

In summary, incorporating these enhancements into the CNN architecture can significantly improve the accuracy, robustness, and interpretability of gender and age prediction in facial recognition systems, ultimately enhancing their practical utility in real-world applications. The flowchart of model used, and the process is shown in Fig. 1.

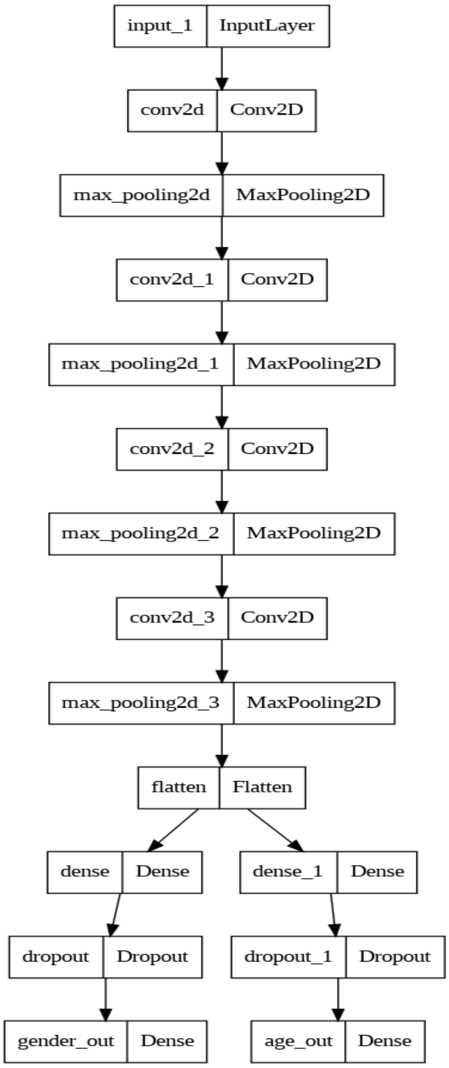

Fig. 1. The flowchart of CNN_MOCM

2.4. Materials

Large datasets of labelled facial images are used to train machine learning models for gender and age detection. CNNs algorithm analyse the visual characteristics of the face, such as facial landmarks, textures, and shapes, to capture important information for gender and age classification. Various software libraries and frameworks provide tools and functionalities for building gender and age detection systems, such as OpenCV, TensorFlow, PyTorch, and sci-kit-learn. These libraries offer pre-trained models, as well as APIs for training custom models and deploying them in applications. Depending on the complexity of the machine learning models used for gender and age detection, training and inference may require significant computational resources. High-performance GPUs (Graphics Processing Units) are often utilized to accelerate training and inference tasks, enabling faster processing of large volumes of facial images. Annotated datasets with labels for gender (male/female) and age groups are essential materials for training gender and age detection models. These annotations provide ground truth labels that the models learn to predict during the training process.

By leveraging these materials, researchers and developers can build accurate and robust gender and age detection systems that can be applied in various real-world applications, such as demographic analysis, targeted advertising, and personalized user experiences.

3. Results

Fig. 2 displays a close-up of a person's face with a neutral to slightly pouting expression. The individual is wearing a yellow headband, and their facial features show a serious or focused demeanour.

Fig. 2. Sample Input Image

Fig. 3 displays a probability density graph of age after applying specific transformations to group the age data into discrete categories. The x-axis represents these age groups, where age 10 is mapped to 0.0, age 20 is mapped to 1.0, ages 30 and 40 are combined and mapped to 2.0, and ages 50, 60, 70, 80, 90, 100, and 110 are grouped together and mapped to 3.0. The graph shows distinct peaks at these points, indicating higher densities at 0.0, 1.0, 2.0, and 3.0, which correspond to the specific age categories mentioned. The presence of these peaks suggests that the majority of the dataset is concentrated within these age groups, with the highest densities observed around these transformed age values.

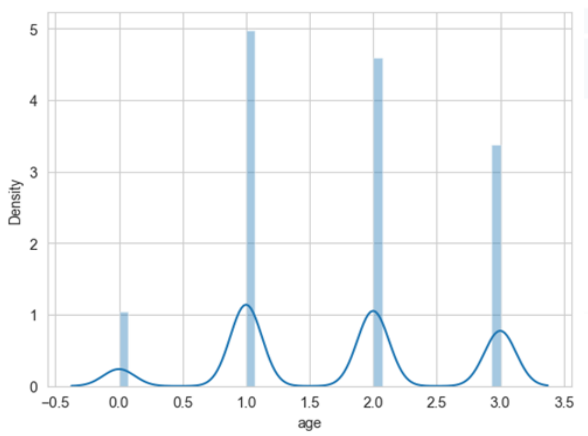

Fig. 3: Probability Density Graph of Age

Fig. 4 shows a probability density graph representing gender distribution, where the x-axis reflects gender values and the y-axis shows density. The gender data has been transformed into two distinct categories: 0.0 likely represents one gender (e.g., female) and 1.0 represents the other (e.g., male). The graph features two prominent peaks at these points, with the first peak at 0.0 and the second at 1.0. Both peaks have densities around 3.0, with the corresponding bar heights reaching approximately 7.0, indicating a significant number of data points in each category. This suggests a fairly balanced dataset between the two gender categories.

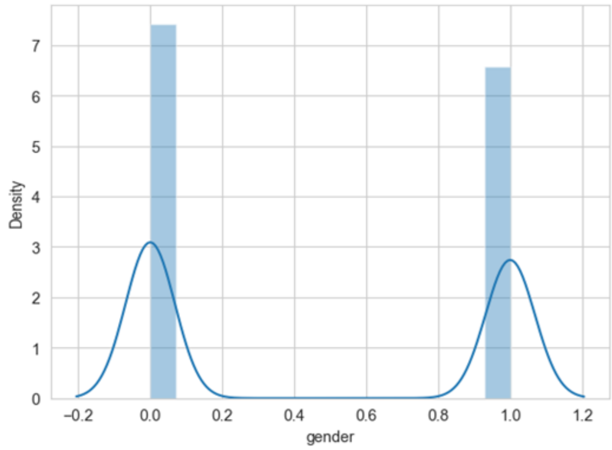

Fig. 4. Probability Density Graph of Gender

The Precision-Recall curve for gender classification in Fig. 5 shows a steady decline in precision as recall increases. Initially, precision remains near 1.0 for lower recall values. However, after reaching a recall value of approximately 0.7, precision begins to sharply decrease, reaching around 0.6 when recall is 0.9 and eventually dropping close to 0.5 as recall approaches 1.0. This indicates that as the model aims to correctly identify more samples (higher recall), the proportion of correctly classified positive samples (precision) decreases. In the ROC curve, the model shows strong performance, with a high true positive rate (TPR) even at low false positive rates (FPR). The curve exhibits a rapid increase with TPR approaching 0.9 at an FPR of about 0.3, indicating a good balance between sensitivity and specificity. The curve gradually flattens as the FPR increases, with the area under the curve (AUC) being close to 1.0, reflecting the model's effectiveness in gender classification across various threshold values.

The Precision-Recall curve for age classification in Fig. 6 exhibits a significant drop in precision as recall increases. Initially, precision starts at 1.0, but it steadily decreases, reaching around 0.8 by the time recall is 0.4. Beyond this point, precision continues to drop, falling below 0.7 as recall exceeds 0.6, and eventually approaches a value close to 0.5 when recall nears 1.0. This pattern indicates a clear trade-off, where the model's ability to correctly identify more age samples comes at the cost of reduced precision. The ROC curve displays a fairly linear relationship between the TPR and FPR. Initially, TPR rises at a moderate pace, reaching 0.6 when the FPR is approximately 0.4. However, unlike an ideal ROC curve, this one does not show a sharp rise in TPR; instead, it follows a more gradual linear path, suggesting a moderate performance in distinguishing age groups. The AUC is lower compared to the gender classification model, indicating that the age classification task may be more challenging for the model.

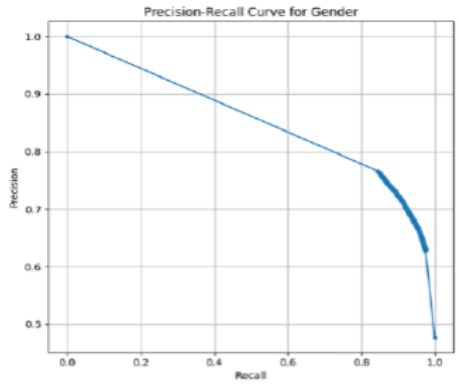

(a)

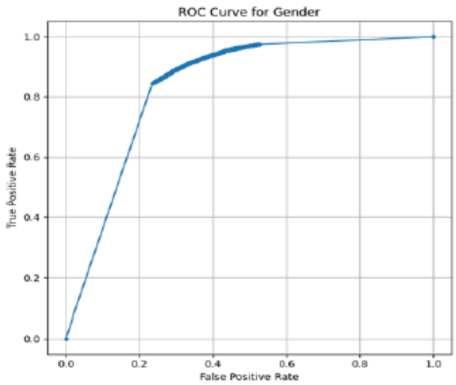

(b)

Fig. 5. Precision – Recall and ROC Curve for Gender using CNN_MOCM

In the gender classification matrix shown in Fig. 7, the model correctly predicted 8,849 true negatives (class "0") and 9,966 true positives (class "1"). However, it made 3,543 false positive predictions, incorrectly classifying samples from class "0" as class "1," and 1,352 false negative predictions, where class "1" samples were incorrectly classified as class "0." This indicates a fairly balanced performance in gender classification, though there are noticeable misclassifications, especially more false positives for class "0." On the other hand, the age classification matrix reveals that the model accurately predicted 4,989 true negatives (class "0") and 18,721 true positives (class "1"), with no false positive or false negative predictions. This suggests that the model performs almost perfectly for age classification, particularly for class "1." However, the absence of false positives or negatives may also point to class imbalance, which could affect the model's generalisation ability in more diverse datasets.

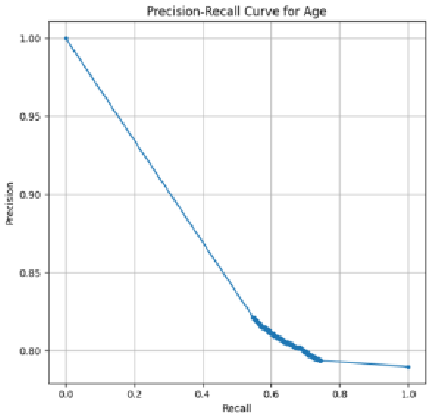

(a)

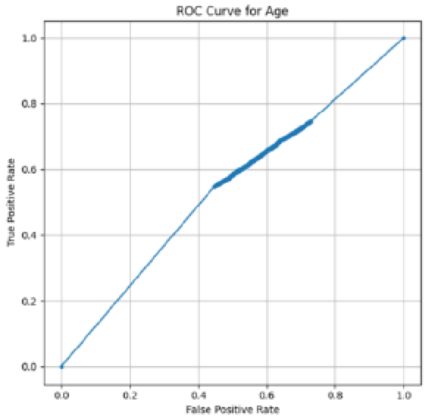

(b)

Fig. 6. Precision – Recall and ROC Curve for Age using CNN_MOCM

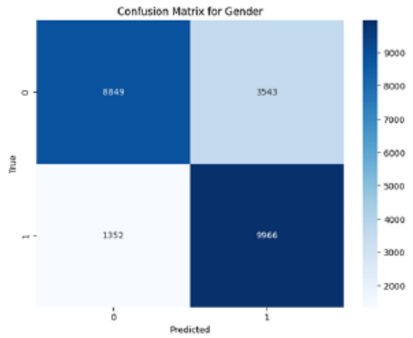

(a)

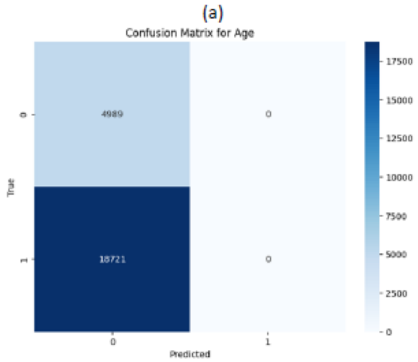

(b)

Fig. 7. Confusion Matrix of Gender and Age using CNN_MOCM

Table 1 showcases a comparative analysis of five algorithms for gender and age detection, highlighting their accuracy. The CNN_MOCM achieves the highest gender detection accuracy at 93.2% and the lowest mean absolute error (MAE) for age prediction, with a maximum deviation of 2 years (i.e., for a 30-year-old, the predicted age could range between 28 and 32 years). For gender detection, the other methods demonstrate accuracy levels as follows: ResNet-50 with Multi-Task Learning achieves 92%, MobileNetV2 with EfficientNet Heads reaches 91%, InceptionV3 with Attention Mechanism records 90%, and VGG-Face with Siamese Network shows 89%. In terms of age prediction accuracy, the CNN_MOCM clearly outperforms the rest with the smallest error margin. Other algorithms have a larger deviation, with ResNet-50 showing a deviation of ±4.5 years, MobileNetV2 with ±4.8 years, InceptionV3 with ±5 years, and VGG-Face demonstrating the largest error margin of ±5.2 years. Overall, the CNN_MOCM surpasses the alternative approaches by providing the most accurate predictions for both gender and age detection, making it a highly effective solution for facial analysis tasks.

Table 1. Quantitative Analysis of Accuracy between 5 different algorithms.

METHODS / ALGORITHMS	DETECTION	ACCURACY (in % and years)
ResNet-50 with Multi-Task Learning	GENDER	92
	AGE	4.5
VGG-Face with Siamese Network	GENDER	89
	AGE	5.2
MobileNetV2 with EfficientNet Heads	GENDER	91
	AGE	4.8
InceptionV3 with Attention Mechanism	GENDER	90
	AGE	5.0
CNN_MOCM (Proposed Method)	GENDER	93.2
	AGE	2

4. Conclusion:

Using CNN_MOCM for gender and age prediction from facial images is promising. CNNs excel at learning intricate facial features crucial for distinguishing gender and age. Challenges include variations in lighting, expressions, and ethnic diversity affecting prediction accuracy, and biases in training data leading to skewed results. To enhance performance, researchers must ensure diverse and representative training data, employ techniques like data augmentation and regularization, and validate models rigorously. Ethical considerations regarding biases and fairness are paramount. Despite progress, ongoing research is vital to improve accuracy, mitigate biases, and address ethical concerns for more reliable and equitable gender and age prediction using CNNs.

References

1. Levi, G., & Hassner, T. (2015). Age and gender classification using convolutional neural networks. In Proceedings of the IEEE conference on computer vision and pattern recognition workshops (pp. 34-42).
2. Xing, J., Li, K., Hu, W., Yuan, C., & Ling, H. (2017). Diagnosing deep learning models for high accuracy age estimation from a single image. Pattern Recognition, 66, 106-116.
3. Eidinger, E., Enbar, R., & Hassner, T. (2014). Age and gender estimation of unfiltered faces. IEEE Transactions on information forensics and security, 9(12), 2170-2179.
4. Kelliher, E. (2021). Human Age and Gender Classification using Convolutional Neural Networks.
5. Lee, B., Gilani, S. Z., Hassan, G. M., & Mian, A. (2019, December). Facial gender classification—analysis using convolutional neural networks. In 2019 Digital Image Computing: Techniques and Applications (DICTA) (pp. 1-8). IEEE.
6. Liu, W., Chen, L., & Chen, Y. (2018, October). Age classification using convolutional neural networks with the multi-class focal loss. In IOP conference series: materials science and engineering (Vol. 428, p. 012043). IOP Publishing.
7. Sundararajan, K., & Woodard, D. L. (2018). Deep learning for biometrics: A survey. ACM Computing Surveys (CSUR), 51(3), 1-34.
8. Hamdi, S., & Moussaoui, A. (2020, December). Comparative study between machine and deep learning methods for age, gender and ethnicity identification. In 2020 4th International Symposium on Informatics and its Applications (ISIA) (pp. 1-6). IEEE.
9. Sharma, N., Sharma, R., & Jindal, N. (2022). Face-based age and gender estimation using improved convolutional neural network approach. Wireless Personal Communications, 124(4), 3035-3054.
10. Wang, X., Guo, R., & Kambhamettu, C. (2015, January). Deeply- learned feature for age estimation. In 2015 IEEE Winter Conference on Applications of Computer Vision (pp. 534-541). IEEE.
11. Meinedo, H., & Trancoso, I. (2011). Age and gender detection in the I- DASH project. ACM Transactions on Speech and Language Processing (TSLP), 7(4), 1-16.
12. Saxena, A., Singh, P., & Singh, S. N. (2021, January). Gender and age detection using deep learning. In 2021 11th International Conference on Cloud Computing, Data Science & Engineering (Confluence) (pp. 719-724). IEEE.
13. Shaker, S. H., & Al-Khalidi, F. Q. (2022). Human Gender and Age Detection Based on Attributes of Face. International Journal of Interactive Mobile Technologies, 16(10).
14. Nada, A. A., Alajrami, E., Al-Saqqa, A. A., & Abu-Naser, S. S. (2020). Age and gender prediction and validation through single user images using CNN. Int. J. Acad. Eng. Res.(IJAER), 4, 21-24.
15. Ghosh, S., & Bandyopadhyay, S. K. (2015). Gender classification and age detection based on human facial features using multi-class SVM. British Journal of Applied Science & Technology, 10(4), 1-15.
16. Brandao, M. (2019). Age and gender bias in pedestrian detection algorithms. arXiv preprint arXiv:1906.10490.
17. Kumbhar, U., & Shingare, A. S. (2021). Gender and Age Detection using Deep Learning. IJSR.
18. Wang, X., Ali, A. M., & Angelov, P. (2017, June). Gender and age classification of human faces for automatic detection of anomalous human behaviour. In 2017 3rd IEEE International Conference on Cybernetics (CYBCONF) (pp. 1-6). IEEE.
19. Karahan, M., Lacinkaya, F., Erdonmez, K., Eminagaoglu, E. D., & Kasnakoglu, C. (2022). Age and gender classification from facial features and object detection with machine learning. Journal of fuzzy extension and applications, 3(3), 219-230.
20. Zaman, S. R., Sadekeen, D., Alfaz, M. A., & Shahriyar, R. (2021, July). One source to detect them all: gender, age, and emotion detection from voice. In 2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC) (pp. 338-343). IEEE.
21. Zhang, Z., Song, Y., & Qi, H. (2017). Age progression/regression by conditional adversarial autoencoder. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 5810-5818).IEEE.
22. Z. Zhang, “UTKFace | Large Scale Face Dataset,” UTKFace | Large Scale Face Dataset. Accessed: Oct. 12, 2024. [Online].

Authors Introduction

Mr. Pratham Gupta

He is currently pursuing his B.Tech degree in Computer Science from School of Computer Science and Engineering, Vellore Institute of Technology, Chennai, India.

Dr. Dhanush R.

Dr. Dhanush R is an Associate Professor in the School of Electronics Engineering, VIT Chennai. His research interests include Neuroscience, Biomechanics. Motor learning, Motor Control.

Dr. Amutha S

Dr. Amutha S is currently working as an Assistant Professor, School of Computer Science and Engineering, Vellore Institute of Technology, Chennai, India. Her area of research includes Data mining, Machine learning, Pattern detection, HCI, IoT/IoMT and Deep learning.

Dr. Heshalini Rajagopal

She received her PhD and Master's degree from the University of Malaya, Malaysia in 2021 and 2016, respectively. She is an Assistant Professor in the School of Engineering and Computing, Mila University. Her research interests include image processing, artificial intelligence and machine learning.