Deep Learning Model for Gender and Age Detection

P.B.Kasalkar[1], S.A.Kumbhar[2]

[1]Student, D.Y.Patil Agriculture and Technical University
[2]Assistant Professor, D.Y.Patil Agriculture and Technical University

Abstract
Age and gender, two significant facial attributes, wield considerable influence within society. The automation of age and gender recognition holds substantial promise across a spectrum of real-world applications, including customer service, priority voting systems, medical diagnosis, and human-computer interaction. Leveraging deep learning techniques has emerged as a common thread in many research endeavours, yielding noteworthy performance enhancements. The integration of diverse deep learning models, coupled with an assessment of accuracy improvements, paves the way for further exploration. The central objective of this paper is to meticulously scrutinize age and gender recognition across various datasets and deep learning models. The paper expounds upon the advancements achieved in this domain, accentuating the contributions made. It outlines the models and datasets employed and rigorously evaluates the proposed approach through the lens of obtained results.

Keywords: Age and gender classification, Deep learning, Deep neural network, Datasets.

I.INTRODUCTION

Age and gender classification involves the identification of an individual's age and gender based on images or videos. Both age and gender are pivotal in person identification. In the modern era of social media, the automatic classification of age and gender via facial images is garnering substantial interest. These attributes represent fundamental facets of facial recognition in social interactions. The human face encompasses traits that define identity, age, gender, emotions, and ethnicity. Consequently, the assessment of age and gender holds significant relevance across numerous applications. Real-world scenarios encompass visual surveillance, electronic customer analysis, crowd behavior assessment, online advertising, product recommendation, law enforcement, restricting minors from purchasing prohibited substances, safeguarding children from harmful online content, forensic investigations, anti-aging treatments, beauty product development, movie casting, and more.

Estimating age with the naked eye poses challenges, given the transformation of facial features over time due to factors like skin changes during middle to old age and growth during adolescence. Age and gender identification, therefore, presents an enduring challenge for researchers, compounded by several common issues.

Computer vision emerges as a solution to these quandaries. An aspect of artificial intelligence (AI), computer vision empowers computers and systems to extract meaningful insights from digital images, videos, and visual inputs, subsequently facilitating informed decisions and recommendations. Thus, the need for an efficient model for age and gender estimation tasks becomes imperative.

Past research often utilized individually crafted features alongside machine learning models. Machine learning involves the study of algorithms that can enhance their performance through experience and data. While these approaches have proven effective on certain datasets, they exhibit deficiencies on others, yielding inconsistent results. Consequently, there has been a shift towards convolutional neural network (CNN)-based deep learning models, capable of training on extensive data, employing intricate algorithms, and autonomously extracting features from facial images. Some noteworthy neural network-based deep learning models include VGG16, VGG19, Xception, ResNet50V2, ResNet101V2, InceptionV3, Inception, ResNetV2, MobileNet, MobileNetV2, DenseNet121, DenseNet169, DenseNet201, NASNetMobile, EfficientNetB0, EfficientNetB1, EfficientNetB2, EfficientNetB3, EfficientNetB4, and more.

Deep learning, a subset of machine learning, closely mimics human learning processes, streamlining and expediting the entire process. Facial age and gender recognition stands out as one of the prime applications of deep learning. Convolutional neural networks (CNNs) are instrumental in this domain, leveraging one or more convolutional layers alongside specialized functions. CNNs excel in automating feature detection, reducing reliance on pre-processing, minimizing human intervention without compromising quality, and optimizing parameter counts.

However, certain drawbacks persist in age and gender recognition, including dataset quality affecting performance, susceptibility of machine learning to mislabeled or noisy data, limitations of traditional classification algorithms in capturing intricate nonlinear relationships within image data, suboptimal feature extraction efficiency and accuracy in deep neural networks, sensitivity to minor alignment changes, misclassification challenges, and complications arising from occlusion, pose, illumination, resolution, and facial expressions in images.

The subsequent sections of this study are structured as follows: Section II provides a literature survey on age and gender recognition approaches, while Section III delves into details concerning previous research model performances and comparisons. Finally, Section IV presents the study's conclusions.

II. LITERATURE SURVEY

In 2017, Zang et al. [2] introduced the concept of a Residual network of Residual network (ROR) for automatic age and gender prediction from facial images under unconstrained conditions. This architecture specifically addresses high-resolution facial image age and gender classification. Two mechanisms were employed to enhance age estimation performance: pretraining by gender and training with a weighted loss layer. ROR was initially pretrained on ImageNet, then fine-tuned on the IMDB-WIKI 101 dataset to further learn face image features, and ultimately tuned on the Adience dataset. This approach achieved high accuracy for gender classification, performing well with high-resolution facial images. The achieved accuracy included 67.37% for age estimation and 93.27% for gender estimation. However, lower age estimation accuracy and the relatively slower speed of the ROR model posed challenges. Age estimation accuracy could sometimes be affected by minor alignment changes, and issues with the Adience dataset were also noted, where the model excelled in specific features only.

In 2018, Philip Smith et al. [3] addressed age and gender recognition through transfer learning using deep CNNs. They employed transfer learning with pretrained VGG19

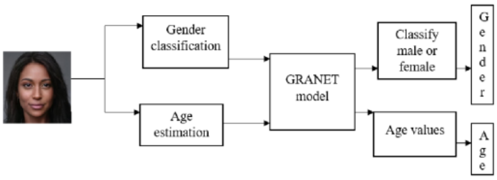

Fig.2. Block diagram of GRANET model, Avishek Garain, et al. [1]

and VGGFace models to enhance efficiency. The study explored training techniques, including input standardization, data augmentation, and label distribution age encoding. The MORPII dataset was utilized, with VGGFace outperforming VGG19. VGGFace demonstrated higher accuracy with a gender prediction rate of 98.68% within a margin of 4.1 years. However, age recognition produced a mean absolute error (MAE) of 4.1 years due to female-specific characteristics. Gender prediction was notably influenced by features like long hair presence and absence, as well as head tilt. The study highlighted the need for larger, cleaner datasets, as mislabeled data and excessive noise could lead to incorrect model learning. Challenges attributed to the MORP II dataset were also acknowledged, including noisy variations such as heads tilted in different directions and an overrepresentation of male, black-colored images.
In 2019, Ningning Yu et al. [4] proposed an ensemble learning approach for facial age estimation using non-ideal facial imagery, as depicted in Fig. [1]. The methodology comprised image preprocessing, feature extraction, and age prediction. The input face image underwent separate preprocessing in RGB Stream, Luminance Modified Stream, and YIQ Stream. Three pretrained deep convolutional neural networks (DCNNs) with softmax were used for feature extraction and age estimation as weak classifiers. The ensemble learning module integrated the outputs of the three weak classifiers to yield a more accurate estimation. The IMDB-WIKI dataset from Wikipedia was utilized. The three-stream method notably improved performance, and ensemble learning enhanced accuracy by fusing the weak classifiers. Evaluation indexes, including Exact Match (AEM) and an error within one age category (AEO), indicated that the ensemble method achieved AEM of 45.57% and AEO of 88.20%. Challenges noted included the need for an efficient global search method and an algorithm that encompasses a more comprehensive dataset. DCNN feature extraction and training times were substantial, and dataset-related issues, particularly ambient illumination and complex backgrounds in IMDB-WIKI, affected estimation difficulty and overall performance.

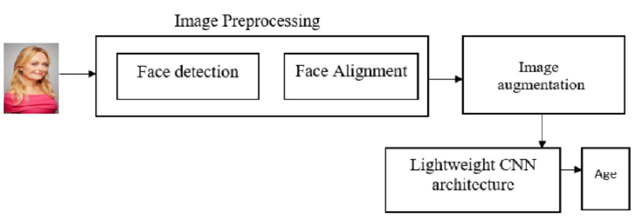

Fig.1. Block diagram for Lightweight CNN Age estimation, Olatunbosun, et al. [5].

In 2020, Olatunbosun et al. [5] introduced a Lightweight Convolutional Neural Network (CNN) for the estimation of real and apparent ages of human faces, as illustrated in Fig [1]. The task of real and apparent age estimation holds numerous real-world applications, including medical diagnosis, forensics, and facial beauty product development. Traditional CNN models often exhibit challenges such as excessive complexity, a high number of network parameters and layers, prolonged training times, and the need for extensive training datasets, leading to increased computation costs and storage overheads. To address these issues, the authors proposed a lightweight CNN architecture with fewer layers, specifically tailored for real and apparent age estimation.

The proposed methodology commences with image preprocessing, encompassing crucial steps such as face detection and alignment. Subsequently, image augmentation techniques are applied, involving random scaling, random horizontal flipping, color channel shifting, standard color jittering, random rotation, and the generation of alternate copies for each training image. The core of the approach lies in estimating real and apparent age utilizing the lightweight CNN model.

The FG-NET, MORP II, and APPA-REAL datasets were employed for evaluation. Notably, FGNET produced a Mean Absolute Error (MAE) of 3.05, MORP II yielded an MAE of 2.31, and APPA-REAL resulted in an MAE of 4.94. The model's performance was notably enhanced when trained on the MORP II dataset. The lightweight CNN exhibited significantly reduced training times, a characteristic that underscores its efficiency.

However, the study highlights the need for more robust and higher-quality image processing algorithms, as the ability to detect unfiltered images quickly is imperative. The emphasis on a lighter CNN architecture with fewer parameters remains crucial. It's noteworthy that age classification challenges posed by non-frontal face images have been mitigated, as they diminish the complexities associated with unfiltered faces.

The study underscores the challenging nature of dataset-related issues, particularly the presence of noisy variations, which significantly impact performance. In conclusion, the proposed lightweight CNN approach demonstrates promise in real and apparent age estimation, offering enhanced efficiency and effectiveness across diverse datasets and scenarios.

In 2021, Garain et al. [1] introduced a model named GRANET (Gated Residual Attention Network) for the classification of age and gender from facial images, as depicted in Fig [2]. Previous research efforts in this domain have revealed common shortcomings, including elevated Mean Absolute Error (MAE), suboptimal age estimation accuracy, and limited gender classification performance. Additionally, model performance has been shown to be susceptible to minor changes in alignment, and variations in resolution impact certain models while not affecting others. To address these limitations, the authors proposed the integration of multiple attention blocks through gating mechanisms to form the GRANET model.

The research leveraged several datasets, including FGNET, AFAD, Wikipedia, UTKFace, and AdienceDB. Among these, UTKFace demonstrated superior performance, achieving a remarkable gender recognition accuracy of 99.2% and an age recognition accuracy of 93.7%. AFAD exhibited an age estimation result with a MAE of 3.10, FGNET recorded a MAE of 3.23, Wikipedia yielded a MAE of 5.45, UTKFace showcased an impressive MAE of 1.07, and AdienceDB yielded a MAE of 10.57. However, AdienceDB achieved a 65.1% age estimation accuracy and an 81.4% gender classification rate.

While the GRANET model exhibited promising results, certain drawbacks were identified. Notably, the classification of images of children posed significant challenges. Instances of misclassification and occurrences of erroneous predictions were also noted. Moreover, the model's performance faced constraints when dealing with obstructed or partially viewed images, indicating a need for enhanced intelligence in such scenarios.

An overarching concern in this research pertains to the dataset issue, which notably impacts performance outcomes. Among the datasets, UTKFace proved superior due to its ability to mitigate the effects of variability in factors like illumination, pose, occlusion, and resolution. By leveraging the strengths of UTKFace, the GRANET model showcases a notable improvement in age and gender classification accuracy, offering promise for further advancements in the field.

A. AGE AND GENDER RECOGNITION TECHNOLOGY

Deep Neural Networks and Transfer Learning: Deep neural networks mimic the structure of the human brain, consisting of interconnected nodes. They typically comprise an input layer, hidden layers, and an output layer. When images serve as input, the nodes in the hidden layers multiply inputs by random weights, calculate results, and pass them to the output layer. Notable algorithms in deep learning include Convolutional Neural Networks (CNNs), Long Short-Term Memory Networks (LSTMs), and Recurrent Neural Networks (RNNs). CNNs, for instance, encompass multiple layers that process and extract features from data. The convolution layer employs various filters for convolution operations, followed by a ReLU layer for element-wise operations. A pooled layer then performs down-sampling, reducing feature map dimensions. Finally, a fully connected layer classifies and identifies images after flattening the matrix from the pooling layer. LSTMs, a type of RNN, specialize in learning and memorizing long-term dependencies. RNNs, on the other hand, facilitate feeding LSTM outputs as inputs for the current phase. Transfer Learning, a machine learning technique, involves reusing a pre-trained model as a starting point for a new task. High-accuracy transfer learning models include VGG, Inception, Xception, and ResNet.

B. DATASETS USED IN AGE AND GENDER RECOGNITION

The training testing of model is very important factor in deep learning. Datasets play a major role in training and testing. The availability of different age and gender recognition datasets are great support for researchers. Each of the datasets has different features too.

MORP II Dataset: This facial age estimation dataset comprises 55,134 facial images featuring 13,617 subjects, spanning ages from 16 to 77 years old. Among these images, approximately 84.6% depict males, and 77.2% belong to individuals of black ethnicity.

ImageNet: Organized based on the WordNet hierarchy, ImageNet is an extensive image database encompassing over 20,000 categories, ranging from everyday items like strawberries and balloons to various other objects.

FG-NET Aging Dataset: This dataset comprises 1002 images featuring 82 distinct subjects, encompassing age ranges from infancy to 69 years old. The images are sourced from personal photograph collections, introducing certain challenges. Variability in image quality, arising from factors like photographic paper quality, illumination, resolution, viewpoint, expression, and occlusions (such as facial hair, spectacles, and hats), can impact the dataset.

AFAD Dataset: The Asian Face Age Dataset serves as an evaluation benchmark for diverse age estimation methods. With over 160,000 facial images accompanied by corresponding age labels, this dataset is tailored for age estimation in Asian faces. Notably, labeled samples within the AFAD dataset span ages from 15 to 40.

Wikipedia Age Dataset: This publicly accessible dataset comprises facial images of various celebrities. To ensure age relevance, images lacking age information due to missing photo dates were excluded. In total, the dataset encompasses 62,328 face images sourced from 20,284 celebrities listed on Wikipedia.

UTKFace Dataset: This extensive dataset offers a broad spectrum of facial images, spanning an age range from 0 to 116. Featuring over 20,000 face images, UTKFace dataset includes labels for age, gender, and ethnicity. The images showcase significant variation in illumination, pose, occlusion, facial expression, and resolution. This dataset serves as a valuable resource for tasks such as face detection, age estimation, age progression or regression, and landmark localization.

Fig.3.Sample images of dataset

ADIENCE DATASET: The Adience dataset comprises photographs captured using digital cameras from smartphones or tablets. The images within the dataset capture significant diversities, including pronounced blurriness (low resolution), occlusions, out-of-plane viewpoints, pose variations, and facial expressions.

BEFA: The Bias Estimation in Face Analytics (BEFA) dataset comprises a collection of 13,431 test images that depict various demographic attributes. These attributes include age categories such as infants, children, teenagers, young adults, middle-aged individuals, and seniors. Additionally, the dataset covers gender distinctions with male and female classifications, as well as ethnic diversity encompassing white, black, Asian, and Indian populations.

III. COMPARISON AND DISCUSSION OF PREVIOUS RESEARCHES

Table 1: Comparative study of age and gender recognition based on previous researches.

Model	**Datasets**	**Advantages**	**Disadvantages**
CNN [13]	Adiences dataset	Attaining a commendable precision of 86.8%, the gender estimation task demonstrates remarkable accuracy. This accomplishment is attributed to a reduction in overfitting risk, facilitated by a decreased parameter count."	The age recognition accuracy for younger individuals is notably lower, registering at 50.7%, primarily attributed to the straightforward nature of its design.
ROR [12]	IMBD-WIKI dataset, ImageNet dataset.	Remarkable accuracy of 93.24% has been attained in the field of gender classification.	The accuracy of detecting lower age groups is recorded at 67.34%

		The model performs effectively when applied to high-resolution images.	
CNN [11]	WEAFD dataset	The accuracy of gender classification is notable, reaching 88%. The dataset consists of labeled facial images intended for classification purposes.	The accuracy for age detection is notably deficient, measuring only 38%.
MTCNN [7]	UTKFace dataset, BEFA dataset	It achieves a significant level of accuracy in gender classification. Specifically, for the UTKFace dataset, the accuracy is 98.23%, and for the BEFA dataset, it reaches 93.72%.	Owing to the restricted range of facial attributes, the accuracy for the age classification task is diminished. Specifically, for the UTKFace dataset, the accuracy stands at 70.1%, while for the BEFA dataset, it records a slightly higher rate of 71.83%
VGG19, VGGFace [3]	MORP-II dataset	VGGFace outperforms. It boasts a remarkable accuracy in gender classification, specifically reaching 98.7%.	The process of age estimation yields a higher Mean Absolute Error (MAE) of 4.1 years. Even a minor alteration significantly impacts the predictive task.

At present, numerous existing research models focus on age and gender recognition. While comparing their performance with early studies, gender recognition demonstrates commendable results, but age estimation presents challenges. When evaluating various models using the UTKFace Dataset, the Facenet model achieves a gender recognition accuracy of 91.2% and an age estimation accuracy of 56.9%. Subsequently, the Finetuned Facenet (FFNet) model yields a gender recognition accuracy of 96.1% and an age estimation accuracy of 64%. The Multitask Cascaded Convolutional Neural Networks (MTCNN) described in [7] achieve remarkable accuracy rates for gender and age, with values of 98.23% and 70.1%, respectively. The Residual Attention Network (RAN) model, as detailed in [6], attains accuracy rates of 97.5% for gender recognition and 85.4% for age estimation. When evaluated on the AFAD dataset, the RAN model achieves an age estimation Mean Absolute Error (MAE) of 3.42, while on the FG-NET dataset, the MAE is 4.05. In the most recent research, the GRANET (Gated Residual Attention Network) model introduced in [1] is applied to five publicly available datasets. The resulting MAE values for FG-NET, AFAD, Wikipedia, UTKFace, and Adience are 3.23, 3.10, 5.45, 1.07, and 10.57, respectively. Notably, UTKFace attains superior gender recognition accuracy at 99.2% and age estimation accuracy of 93.7%. Refer to Table 1 for a comprehensive comparative analysis of previous research studies.

Implementing an Age and Gender Detection system using Mask R-CNN involves adapting the Mask R-CNN architecture to simultaneously perform instance segmentation and age/gender classification.

International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

1. **Dataset Preparation:** Gather a dataset of images annotated with instance masks, age labels, and gender labels. The dataset should include images of people with instance masks outlining their bodies and annotations indicating their age (e.g., young, middle-aged, old) and gender (male, female).
2. **Model Architecture:** Extend the Mask R-CNN architecture to accommodate the age and gender classification tasks. The architecture consists of three main components:
   a. **Backbone**: Choose a suitable backbone architecture, such as ResNet, ResNeXt, or another convolutional neural network. The backbone extracts features from the input images.
   b. **Region Proposal Network (RPN)**: This generates region proposals (candidate object bounding boxes) based on the features extracted by the backbone.
   c. **ROI Align and Classification Heads**: Adapt the ROI Align and classification heads of Mask R-CNN. Add separate classification heads for age and gender. These heads take the region of interest (ROI) features from the ROIs proposed by the RPN and perform age and gender classification.
3. **Loss Function:** Define a multi-task loss function that combines the losses from instance segmentation, age classification, and gender classification. The overall loss is a weighted sum of these individual losses. Adjust the loss weights based on the importance of each task.
4. **Training:** Train the model using the annotated dataset. During training, the model will simultaneously learn to segment instances and classify age and gender. Utilize a training loop that updates the model's weights using backpropagation and gradient descent. Monitor the loss values and validation metrics to track the model's progress.
5. **Inference and Post-Processing:** During inference, run the trained model on new images. The model will generate instance masks, age predictions, and gender predictions for each detected person. Apply suitable post-processing steps to refine the predictions, such as thresholding, non-maximum suppression for instance segmentation, and selecting the most probable age and gender predictions.
6. **Evaluation:** Evaluate the model's performance using appropriate metrics for age and gender classification accuracy, as well as instance segmentation metrics like mean Average Precision (mAP) or Intersection over Union (IoU).
7. **Fine-Tuning and Optimization:** Based on the evaluation results, fine-tune the model, adjust hyperparameters, or explore different backbone architectures to improve performance.

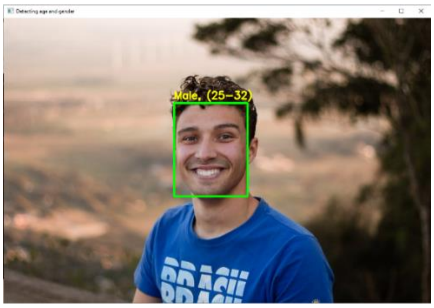

Fig 4. Detection of Age and Gender (Male)

Fig 5. Detection of Age and Gender (Female)

Result and Discussion:

In the experimentation phase, our dataset comprised approximately 30,000 images, and these were utilized to train the model. The proposed model was developed using Python and coded within the Visual Studio platform. Execution of the modules was carried out via command prompt. The successful culmination of training involved the utilization of Caffemodel, a deep learning framework encompassing learning methodologies. This process yielded a ".prototxt" file delineating the neural network's architecture.

Initiating the workflow, we commence with facial detection to identify facial components within individuals. Subsequently, we harness the potency of feature extraction techniques, pivotal in unearthing distinct attributes from specific images. The outcome thereof is seamlessly relayed to the age detection and prediction system, which undertakes the task of prognosticating the corresponding age group.

A pivotal juncture follows as the obtained result is seamlessly channeled to the executive system, which imparts predictions based on cumulative experience. This iterative refinement lends a multi-faceted dimension to age group prognostications. The narrative evolves further as the torch is passed to the gender detection module, meticulously distilling gender determinations and sharing these insightful outcomes with the external realm.

Illustrative instances of output substantiate the methodology's efficacy. In a specific scenario, an image of a woman serves as the input. The subsequent module execution transpires through the following command:

CONCLUSION

Facial age and gender recognition pose a complex challenge, yet its significance in society cannot be understated due to its numerous real-world applications. While attempting age and gender recognition using individually tailored features and traditional machine learning methods, the limitations become apparent. These methods struggle to capture the non-linear relationships present in images. Thus, the emergence of deep learning models has proven highly consequential.

Recent studies have shown that deep learning models outperform their predecessors and contribute to enhanced performance. Gender classification tasks, in particular, have demonstrated notable advancements and favourable outcomes. However, there remains a significant room for improvement in age estimation. Current single-model approaches are limited in their ability to extract diverse features, consequently impacting overall performance.

To address these challenges, the future of age and gender recognition hinges on the introduction of ensemble techniques employing various deep learning models. This approach holds promise for elevating performance levels and achieving substantial advancements in the field.

REFERENCES:

1. A Garain, B Ray," GRANet: A Deep Learning Model for Classification of Age and Gender From Facial Images" IEEE Access, vol.9, pp.85672-85689, 2021.
2. K. Zhang, C. Gao, L. Guo, M. Sun, X. Yuan, T. X. Han, Z. Zhao, and B. Li,``Age group and gender estimation in the wild with deep RoR architecture,"IEEE Access, vol. 5, pp.22492-22503, 2017.
3. P. Smith and C. Chen, ``Transfer learning with deep CNNs for gender recognition and age estimation," in Proc. IEEE Int. Conf. Big Data (Big Data),Dec. 2018, pp. 2564-2571.
4. Ningning yu, l. qian, y. huang, and yuan wu," Ensemble Learning for Facial Age Estimation Within Non-Ideal Imagery" Facial, IEEE Access, vo.7, pp.97938- 97948,2019.
5. Olatunbosun agbo-ajala and Serestina viriri" A Lightweight Convolutional Neural Network for Real and Apparent Age Estimation in Unconstrained Face Images", IEEE Access, vol.8, pp.162800-162808,2020.
6. F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang "Residual attention network for image classification," in Proc. IEEE Conf.Comput. Vis. Pattern Recognit. (CVPR), Jul. 2017, pp. 3156-3164.
7. A. Das, A. Dantcheva, and F. Bremond, ``Mitigating bias in gender, age, and ethnicity classification: A multi-task convolution neural network approach," in Proc. Eur. Conf. Comput. Vis. (ECCV) Workshops, 2018, pp. 1-13.
8. A.V. Savchenko,"Efficient facial representations for age, gender and identity recognition in organizing photo albums using multi-output ConvNet,"PeerJ Comput. Sci., vol. 5, p. e197, Jun. 2019,
9. J. H. Lee, Y. M. Chan, T. Y. Chen, and C. S. Chen, ``Joint estimation of age and gender from unconstrained face images using lightweight multitask CNN for mobile applications," in Proc. IEEE Conf. Multimedia Inf.Process.Retr. (MIPR), Apr. 2018, Art. no. 17877533.
10. S. Hosseini, S. H. Lee, H. J. Kwon, H. I. Koo, and N. I. Cho, ``Age and gender classification using wide convolutional neural network and Gabor filter," in Proc. Int. Workshop Adv. Image Technol., 2018, pp. 1-3,
11. N. Srinivas, H. Atwal, D. C. Rose, G. Mahalingam, K. Ricanek, and D. S. Bolme, "Age, gender, and _ne-grained ethnicity prediction using convolutional neural networks for the East Asian face dataset," in Proc. 12th IEEE Int. Conf. Autom. Face Gesture Recognit. (FG), May 2017,pp. 953-960
12. K. Zhang, C. Gao, L. Guo, M. Sun, X. Yuan, T. X. Han, Z. Zhao,and B. Li, ``Age group and gender estimation in the wild with deep RoR architecture," IEEE Access, vol. 5, pp.22492-22503,2017,
13. G. Levi and T. Hassner, "Age and gender classification using convolutional neural networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit.Workshops, Jun. 2015, pp. 34-42,