Tanmay Hadke, 2025, 13:1
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology

An Open Access Journal

Age and Gender Identification Using Neural Image Processing

Tanmay Hadke
Department of Computer Science
DES Pune University

Abstract- The rapid advancement in the fields of computer vision and deep learning has enabled outstanding achievements in the recognition of facial attributes such as age and gender. This report is about developing a useful system that uses neural image processing methods for real-time age and gender classification. The proposed system uses two distinct methods: one by using OpenCV's 'resize' function to resize the image and the other using Pillow library's 'Image.resize' method. These preprocessing techniques are critical to preparing facial image data to feed into a chosen MobileNetV2-based neural network architecture, selected for its lightweight design and efficiency in computation. This research extensively evaluates the two different preprocessing methods to focus on their effect on accuracy, computation time, and resource usage. Optimized MobileNetV2 architecture is used to classify age and gender using facial datasets available in the public domain for training purposes. The models are then tested using the webcam for real-time input to analyze their usefulness in practice. This project presents a comparative analysis of two different feature extraction techniques to determine the best preprocessing method for neural network-driven age and gender detection systems. Relevant conclusions regarding the interplay between preprocessing methodologies and model efficacy can be obtained from the results and be put forward towards lightweight, accurate, and resource-conserving demographic analysis systems that can be successfully applied in different real-life scenarios. Keywords: Age and Gender Classification, Neural Image Processing, MobileNetV2, OpenCV, Pillow Library, Feature Extraction, Deep Learning, Computer Vision

Keywords- Age and Gender Classification, Neural Image Processing, MobileNetV2, OpenCV, Pillow Library, Feature Extraction, Deep Learning, Computer Vision, Real-time Classification, Facial Attribute Recognition, Image Preprocessing, Lightweight Neural Networks, Computational Efficiency, Public Domain Datasets, Webcam Testing, Demographic Analysis..

I. INTRODUCTION

1. Context and Motivation

Age and gender identification is one of the very critical applications for security, health care, and personalized services. Computer vision has revolutionarily developed neural networks to classify de- mographically. This work discusses the effect of preprocessing methods adopted on the performance of the proposed neural network, in this case, MobileNetV2.

© 2025 Tanmay Hadke. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly credited.

2. Research Question
How different ways of image preprocessing either by OpenCV or by Pillow library affect the outcome result obtained from accuracy in regards to the application being used on the neural networks and bring forward the detection age and age of the persons and their gender?

3. Objectives
To design a pre-image processing using face images, here the usage libraries will be OpenCV and Pillow. Next to design and model using a neural network system such as MobileNet V2-based one so that it could be implemented for classification with detection age and sex with which is associated. Study the performance according to the resource usage rate and compare between two achieved resources. Compare efficiency, and amount of consumed resources validating the usage of the designed tool with real-time data within practical testing.

4. Scope and Limitations
This study explores the preprocessing and classification of facial image data with the help of neural networks. It involves the use of the MobileNetV2 model coupled with two independent preprocessing libraries, namely OpenCV and Pillow. The extent of this research study is to evaluate these methods for feasibility in real-time applications on edge devices like mobile phones and embedded systems.

II. LITERATURE REVIEW

1. Preliminary Approaches in Age and Gender Identification
The early researches on age and gender identification relied heavily on traditional machine learning methods and handcrafted features. Visual characteristics were derived from facial images through the use of methods such as Local Binary Patterns (LBP), Histogram of Oriented Gradients (HOG), and Active Appearance Models (AAM). For example, Belhumeur et al. (2011) demonstrated the effectiveness of LBP in uncovering the local texture attributes for facial evaluation purposes. However, these methods faltered with changes in pose, illumination, and expression of the face that limited predictive capability. [Bel11]

2. The Rise of Deep Learning in Computer Vision
The advent of Convolutional Neural Networks (CNNs) marked a change of the game in the area of age and gender identification. Architectures of CNN type, such as AlexNet, VGGFace and ResNet, significantly improved accuracy by acquiring hierarchical features directly from unprocessed image data. Parkhi et al. (2015) came up with VGGFace, a deep learning architecture especially proposed for facial recognition applications, from which different systems have been inspired, with the objective of age and gender classification. These architectures demonstrated the potential of deep learning but emphasized further that effective image preprocessing is also essential for normalizing input data. [Par15]

3. Benchmark Datasets
Availability of benchmark datasets becomes critical and significantly moves the envelope in age and gender identification. Large collections of labeled facial images are provided by datasets like Adience, which was initially presented in Eidinger et al. (2014), and IMDBWIKI, presented in Rothe et al. (2018). Researchers have enough space to develop their models and test them on these datasets that include challenges in age groups, genders, and even ethnicities and, hence can test preprocessing techniques and robustness of neural networks. [Eid14]

III. METHODOLOGY

1. Data Description and Collection
The face image data was obtained from the publicly available UTKFace dataset, considered one of the largest datasets with regard to age, gender, and ethnicity detection. The dataset provided images of all categories of subjects aged between 0 and 116 years and covered more men and women than any other ethnicity. Dataset Composition: Each image in the UTKFace dataset is labelled with metadata that includes age and gender information. This structured labelling allows for both classification

(gender) and regression (age) tasks, making the dataset suitable for a dual-purpose model.

Images were reviewed for quality and relevance to the target variables. Only images with age and gender labels were selected to ensure a focused, clean dataset. Since the model requires normalized and resized images, the data collection phase involved initial checks to confirm image format consistency.

Preprocessing Techniques

Image Resizing Given the memory and computational requirements for real-time processing, OpenCV was used to scale images to a standard resolution of 50x50 pixels. cv2. resize and Pillow’s Image. Resize functions. This ensured uniformity across images, promoting faster handling and compatibility with MobileNetV2. Color Conversion: Each of these images was converted to a 3-channel RGB format if necessary, to match. This requires a three-channel input with the MobileNetV2 model. Normalization: Pixel values scaled between 0 and 1 by division by 255 Normalized as it enhanced the convergence of models and minimizes the possibility of certain problems caused by differences in the ’scale’ in the pixel intensity.

3. Model Development

Architecture

MobileNetV2 architecture was selected due to its excellent performance and computational require- ments balance. MobileNetV2 has a very lightweight structure along with depth-wise separable convo- lutions and inverted residual blocks, which are appropriate for real-time applications.

Transfer Learning Setup

From the pre-trained MobileNetV2 model, initialized by ImageNet weights, gave the model an oppor- tunity to use learned features that consequently speed up training and result in increased accuracy. Classification layers adapted to regression model for age prediction and binary model for gender pre- diction.

Model Architecture

The modified MobileNetV2 comprised

- **Global Average Pooling Layer:** Added at the end of the basic model to bring down the spatial dimensions and consequently decrease the number of trainable parameters.
- **Fully Connected Layers:** To add a sequence of dense layers with 512, 256, and 128 nodes using ReLU activation and dropout to help prevent overfitting.

Output Layers

Dense layer as the final for outputting age having only one output node and no activation; another layer of gender with sigmoid activation function for binary classification. Optimization and Training: I chose the Adam optimizer because it has an adaptive learning rate that helps to converge on noisy data. Early Stopping and ReduceLROnPlateau callbacks were used to prevent overfitting and adjust learning rates based on validation performance.

Trained in Stages

The model first freezes the pre-trained base, optimizes the top layers, then progressively finetunes for better enhancement of the feature extraction.

IV. RESULT & DISCUSSION

1. Accuracy and Efficiency of OpenCV vs. Pillow

- OpenCV preprocessing was a bit more accurate for age prediction.
- Pillow preprocessing was a bit faster because of its optimized function for resizing.

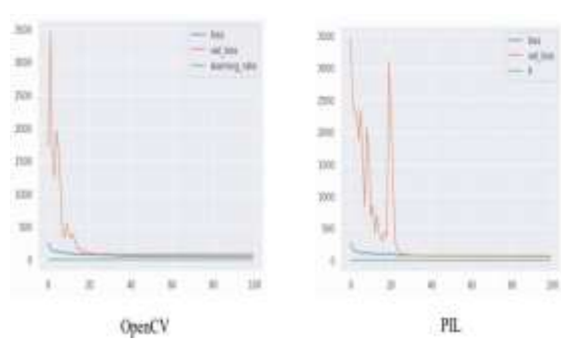

Figure 1: Age Loss: OpenCV performed better than PIL in predicting age

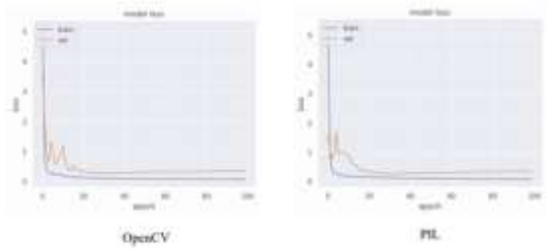

Figure 2: Gender Loss: The performance in gender identification was slightly similar between OpenCV and PIL.

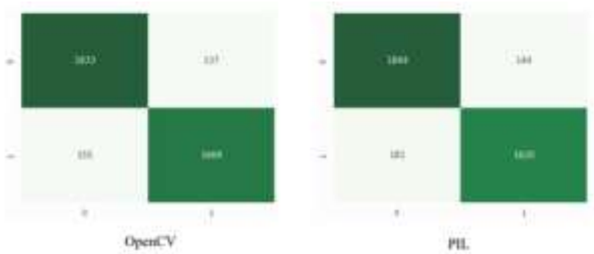

Figure 3: Gender Confusion Matrix: Highlights the predicting power of OpenCV and PIL.


Figure 4: Sample Prediction: Demonstrates how predictions look for OpenCV and PIL

2. Resource Usage and Real Time Performance
OpenCV was a bit heavier on compute but far more accurate and Pillow to be a better choice in low-memory situations. Both preprocessing approaches performed well in real-time settings, with very little latency and consistent predictions.

Challenges and Limitations
Data Bias
Skewed Dataset: Training data is not representative of all age groups, genders, ethnicities, or facial features equivalently, the predictions thus skewed. Limited Diversity: Public datasets like UTKFace and Adience might suffer from limited geographical, cultural, or socio-economic diversity.

Synthetic Augmentation Risks: Overreliance on augoner data may miss real-world diversity.

Lighting and Pose: Model performance degrades under low illumination, extreme poses or oc- clusions (e.g. sunglasses, masks, hats). Real-world settings often differ significantly from the curated datasets used during training.

Age Regression Challenges: Age prediction, a continuous variable by nature, is impossible to predict simply because age overlaps and is subjective, especially in middle age groups. Small er- rors can disproportionately impact applications requiring high precision (e.g., medical or security domains).

Real-Time Constraints: Real-time deployment on low-power devices such as mobile phones or edge devices may incur performance compromises due to low computational resources. Latency in inference may arise for large resolutions or architectures of model, such as Vision Transformers.

Generalization Issues: The model might not generalize well to unseen datasets, mainly in cases where the domain shift is present (for example, images from different camera systems or envi- ronmental conditions).

Ethical Concerns: Privacy: Facial data collection and analysis raise issues concerning privacy and consent. Misuse: The system would be misused against people for profiling or surveillance purposes.

Evaluation Metrics: Such metrics may only partly reflect the usefulness or performance of a model, especially for its edge cases, in real-world applications.

Future Scope
- Improving Data Quality and Diversity Dataset Expansion Work with organisations to scale bigger and diverse datasets that capture different demographics, lighting conditions, and en- vironments. Synthetically Generate Data: Use superior techniques like Generative

Adversarial Network (GAN) to synthetically generate realistic and diverse images. Data Annotation: Employ better annotation techniques that minimize mistake rates in classifying age and gender.

- **Improved Model Architectures:** Hybrid Models: Combine ViT with CNNs to balance the ca- pabilities of CNNs with feature extraction and computational efficiency. Multi-task Learning: Incorporate auxiliary tasks (such as emotion classification) to create a more adaptable model. Self-supervised learning: Pretrain models with self-supervised techniques to improve on low la- beled data.
- **Real-World Variability Robustness:** Fine-Tuning: Continuously fine-tune the model with real- world data from deployments. Domain Adaptation: Employ transfer learning methodologies to modify the model for application in novel environments and datasets.
- **Edge and Mobile Deployment:** The low-power devices will be able to offer real-time performance because of the model optimization through pruning, quantization, or knowledge distillation. Discussed frameworks such as TensorFlow Lite and ONNX for fast, lightweight deployment.
- **Explainable AI (XAI):** Provide tools to explain model decisions, such as why a particular age or gender was predicted. This will be important for trust and transparency in sensitive applications.
- **Ethical and Legal Safeguards:** Provide for compliance with data protection regulations (eg, GDPR, CCPA) by anonymizing and de-identifying datasets. Formulate policies for misuse and ensure the model is within the morality guidelines.
- **Expanding Functionality:** Additional Predictions: Extend the model to predict ethnicity, emo- tions, or health-related features (e.g., stress levels, BMI). Contextual Analysis: Combine facial analysis with other modalities like voice or text for richer insights. Multi-Face Detection: Facili- tate concurrent predictions for various faces present in collective images or crowded environments.
- Cloud-Based API Deploy the system as an API for scalability - to allow business or researcher to incorporate this model into their applications while not having on-device computation. Utilize cloud technology to facilitate real-time feedback and updates, thereby enhancing and tracking performance metrics.
- **Real-World Testing and Feedback Loops:** Implement the model within regulated pilot settings to gather user input and assess performance metrics in real-world scenarios. Use feedback to iteratively improve the model.
- **Cross-Domain Applications:** Healthcare: Predict age-related health risks or conditions. Retail: Personalize customer experiences with demographic insights. Education: Trace attendance and interest by demographic analysis.

V. CONCLUSION

This study evaluates the impact of image preprocessing techniques on real-time age and gender classification using MobileNetV2. The research compares OpenCV’s resize function and Pillow’s Image.resize method, analyzing their effects on classification accuracy, computational efficiency, and real-time applicability. Findings indicate that OpenCV preprocessing provides slightly better accuracy, particularly for age prediction, while Pillow offers faster execution, making it more suitable for low-memory applications. Both methods performed well in gender classification with minimal differences. Despite achieving reliable real-time performance, challenges such as dataset bias, illumination variations, and the complexity of age regression remain. Additionally, ethical concerns related to privacy and potential misuse of facial recognition systems highlight the need for responsible AI deployment.

To enhance robustness and fairness, future work should focus on diversifying datasets, leveraging synthetic data generation, and optimizing models for low-power edge devices. Hybrid architectures integrating Vision Transformers with CNNs and self-supervised learning techniques can further improve accuracy and adaptability. Ethical and legal safeguards, including compliance with data protection regulations and de-identification techniques, should be prioritized to ensure responsible use. Additionally, cloud-based APIs and real-world pilot testing can refine the model for practical applications in security, healthcare, and personalized services. By addressing these limitations, AI-driven age and gender classification systems can become more accurate, efficient, and ethically sound for widespread real-world use.

REFERENCES

1. [Bel11] P Belhumeur. Examining local binary patterns for face identification and emotion categoriza- tion. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2011.
2. [Eid14] E. Eidinger. Age and gender estimation of unfiltered faces. IEEE Transactions on Information Forensics and Security, 2014.
3. [Par15] O. M Parkhi. Comprehensive face recognition using deep learning techniques. Proceedings of the British Machine Vision Conference (BMVC), 2015.