

AGE AND GENDER ESTIMATION IN FACIAL RECOGNITION

Venugopal D R¹, Vinod Kumar², Dr. Mohammed Raffi³

^{1,2}Students, Dept of Computer Science and Engineering, University B D T College of Engineering,
Davangere-577004, Karnataka, India.

³Professor, Dept. of Computer Science and Engineering, University B D T College of Engineering,
Davangere-577004, Karnataka, India.

ABSTRACT

This review paper explores recent advancements in age and gender estimation within the realm of facial recognition technology. The rapid evolution of computer vision and machine learning techniques has led to significant progress in accurately determining both age and gender from facial images. The paper begins by summarizing key findings from existing literature, highlighting the diverse range of methodologies employed, including traditional feature-based approaches and more contemporary deep learning models. A comprehensive analysis of benchmark datasets used for training and evaluation is presented, shedding light on the challenges and limitations associated with different datasets. Furthermore, the review discusses the impact of demographic biases on age and gender estimation models, emphasizing the need for ethical considerations in algorithm development. The paper explores the influence of various factors such as facial expressions, lighting conditions, and cultural variations on the accuracy of predictions. Additionally, it addresses the ongoing efforts to improve the robustness and generalization capabilities of these models across diverse populations.

The synthesis of findings from multiple studies provides a nuanced understanding of the state-of-the-art techniques, their strengths, and potential areas for improvement. The review concludes by proposing future research directions, emphasizing the importance of addressing ethical concerns, mitigating biases, and enhancing interpretability in age and gender estimation models for facial recognition systems. This comprehensive examination serves as a valuable resource for researchers, practitioners, and policymakers in the field of computer vision and facial recognition technology.

Key Words: TensorFlow, Mask R-CNN, Deep Learning, convolutional neural networks.

1. INTRODUCTION

In the contemporary landscape of intelligent systems, the ability to extract meaningful information from facial images has catalyzed groundbreaking advancements across diverse domains. Among the myriad applications, age and gender prediction stand out as pivotal elements influencing social relationships, security protocols, marketing strategies, and personalized user experiences. This review endeavors to present a thorough analysis of existing literature, methodologies, and challenges in the realm of gender classification based on facial recognition and age detection.

The integration of facial recognition technologies into various facets of our digital era, including social understanding, biometrics, identity verification, and crowd behavior analysis, has propelled these technologies to the forefront of innovation. The significance of accurately estimating age and gender from facial images is underscored by its implications in access control, human-computer interaction, marketing intelligence, and visual surveillance. The dynamic nature of these applications necessitates a robust and versatile approach to facial recognition.

The study under consideration leverages deep learning, specifically utilizing the Caffe framework and convolutional neural networks (CNNs), to predict gender and age from single facial captures. Recognizing the challenges posed by intra-class differences, unconstrained photography, and the influence of external factors such as cosmetics and lighting, the authors adopt a classification-based approach for age prediction. Age ranges, delineated in the final SoftMax layer, offer a pragmatic solution to the intricacies of predicting age accurately from a single shot.

The methodology employed integrates the capabilities of the Caffe framework, surpassing traditional benchmarks, and demonstrates superior performance compared to TensorFlow in internal testing. The study's implementation involves the use of a Python script incorporating OpenCV for real-world application. Beyond the technical nuances, this paper delves into the broader implications of age and gender prediction, exploring its impact on social relationships, access control mechanisms, and the burgeoning field of personalized healthcare.

Furthermore, a critical aspect of this review focuses on the existing literature and methodologies employed in gender classification and age detection. By scrutinizing the strengths and limitations of traditional machine learning algorithms, deep learning models, and hybrid techniques, the authors aim to provide a comprehensive understanding of the progress made in these areas. The importance of accurate age detection in personalized healthcare, age-specific marketing, and social robotics is underscored, emphasizing the practical implications of these advancements.

As we navigate through this review, we will delve into the intricacies of datasets typically used for developing and testing gender classification and age detection models. The diversity, size, and representativeness of these datasets will be explored, shedding light on the critical role of objective and inclusive data collection in mitigating algorithmic biases.

In essence, this review seeks to not only contribute to the existing body of knowledge in facial recognition but also serve as a catalyst for future advancements in age and gender identification. By addressing entrenched biases, improving model interpretability, and navigating the complexities of facial variations, we aspire to chart the course for a more robust and equitable future in facial recognition technology.

2. LITERATURE SURVEY

✓ Gender and Age Estimation System from Face Images (SICE 2003 Annual Conference):

- Fukano's work introduces a comprehensive system for gender and age estimation.
- Key elements include face candidate extraction based on skin color and 37 facial foci points.
- Utilization of Gabor Wavelet Transform (GWT) for precise face and parts positioning.
- Facial Age Estimation by Curriculum Learning:
- Wang and team propose a curriculum learning approach for facial age estimation.
- Significantly reduced prediction errors compared to conventional training on the AFAD database.
- Adoption of ResNet-34 CNN architecture for age estimation framework.

✓ Age Prediction Based on Feature Selection:

- Wang, Song, and Liu focus on age prediction using feature selection techniques.
- Ternary classification into age groups, highlighting the impact of selected n-grammes.
- Examination of age-related patterns among Microblog users.

✓ Mask R-CNN:

- He, Gkioxari, Dollár, and Girshick present the Mask R-CNN system for segmenting object instances.
- Outperformance in COCO challenge set tasks, offering enhanced instance segmentation and bounding box detection.
- Age Detection with Face Mask using Deep Learning and FaceMaskNet-9:
- A response to age identification challenges with face masks using FaceMaskNet-9.
- Utilization of a diverse dataset for training, including individuals with and without masks.
- Incorporation of multiple strategies, such as Haar cascade classifier and Gabor filter, for age prediction.

✓ Gender Classification using Face Recognition:

- Bissoon and Viriri address gender classification using PCA and LDA techniques.
- Hybrid approach (FEBFRGAC algorithm) integrating image normalization, feature extraction, and classification.
- Achieved high accuracy in gender classification, particularly with Canny edge detection.

✓ Deep Learning for Gender Recognition:

- Deng, Xu, Wang, and Sun propose a deep learning-based gender detection model.
- Emphasis on a comprehensive dataset including variations in age, race, occupation, and dress style.
- Successful integration of building blocks like ReLU, fine filter, dropout learning, and very deep architecture.

3. METHODOLOGY

Brief overview of the challenges faced in accurate gender and age detection, especially in scenarios involving face masks, prompted by events like the COVID-19 epidemic. Figuring out gender and age is harder when people are wearing masks. There are many types of masks, people wear them differently, and the lighting or background can change a lot. The current methods struggle with these issues, and we need better ways to make them work in real-life situations. Importance of addressing these challenges for applications in pandemic situations, surveillance, and social behavior research. Introduction to the proposed approach using Mask R-CNN to improve Exploration of the effects of mask sorting, obstacles, and lighting conditions on the accuracy of gender and age detection. Discussion of techniques, including information augmentation methods and data linking, to address challenges introduced by these components.

Discussion of potential applications in observational frameworks, public health, and social behavior research. Emphasis on the contribution to the advancement of gender and age detection innovation. Addressing the evolving need for robust and versatile frameworks in the enforcement of social orders. To make this work, we have a plan. We'll collect lots accuracy and robustness in gender and age detection in of pictures of faces with masks, mark them with masked faces.

Overview of existing gender and age detection algorithms and their limitations in mask-dominant scenarios. Introduction to Mask R-CNN as a region- based convolutional neural network designed for object detection and instance segmentation. Importance of combining object detection and semantic segmentation in addressing challenges posed by face masks. Our solution is to use something called Mask R-CNN. It's like a smart computer program that can look at a picture of a face, figure out where the face is, and even understand details like whether the person is wearing a mask or not. It's like teaching a computer to see faces and masks in a smart way.

Explanation of the proposed strategy involving the adaptation of Mask R-CNN for gender and age detection. Detailed steps, including dataset collection, annotation, and the creation of an R-CNN mask to capture gender- and age-specific characteristics.

Iterative preparation and fine-tuning of the model to optimize face recognition and classification performance, addressing enclosure proximity challenges.

Detailed explanation of the dataset creation process, emphasizing diversity in mask styles, poses, lighting scenarios, and ethnicities. Considerations for dataset editing to ensure accurate display and reduce mutilation. Importance of dataset generalizability for model effectiveness.

Discussion of extensive tests using benchmark datasets to assess the proposed approach. Comparison with existing state-of-the-art gender and age detection strategies. Utilization of performance metrics such as accuracy, precision, recall, and F1 scores to evaluate model adequacy. information about gender and age, and then use these pictures to teach our computer program. We'll train it to understand different faces, especially when they're wearing masks. It's like showing the program lots of examples so it gets really good at recognizing faces and figuring out age and gender.

We'll make sure our collection of pictures includes all sorts of faces, with different types of masks and lighting. It's important that our program can handle all kinds of situations, so we need a diverse set of pictures. We'll edit the pictures carefully to make sure they show the right information and won't confuse our program.

After training our program, we'll test it with other pictures to see how well it can figure out gender and age, even when people are wearing masks. We'll compare it with other methods to make sure it's doing a good job.

We'll also look at how our program handles challenges like different types of masks, obstacles in the way, or changes in lighting. We'll use clever techniques to help our program deal with these tricky situations.

Our solution can be useful in many areas like public health and understanding how people behave. By making our program smart in figuring out gender and age despite masks, we contribute to making our communities safer and more orderly.

In summary, our paper introduces a clever way to use Mask R-CNN to solve the puzzle of figuring out gender and age when people are wearing masks. This can be a big help in keeping our communities safe and understanding how people interact, especially in times where masks are a common part of daily life.

Module 1: Data Preparation

Objective:

- ✓ Clean and preprocess photos for compatibility with the Caffe framework.

Tasks:

- ✓ Develop a Python script to handle image preprocessing, ensuring compatibility with Caffe.
- ✓ Clean and preprocess images, considering factors like noise, lighting, and resolution.
- ✓ Save the preprocessed images in a Caffe- compatible format, ensuring consistency in data representation.

Module 2: Model Definition

Objective:

- ✓ Select a CNN architecture and define its settings.

Tasks:

- ✓ Choose a suitable CNN architecture for face recognition (e.g., VGG, ResNet, etc.)
- ✓ Create a `.prototxt` configuration file to define the architecture's structure, including details like layer types, sizes and connectivity.
- ✓ Specify input and output layers, activation functions, and any other relevant settings.

Module 3: Definition of a Solver

Objective:

- ✓ Set up a solver to optimize the chosen model during training.

Tasks:

- ✓ Develop a solver configuration file (e.g.,
`.prototxt`) to define parameters for optimization.
- ✓ Specify solver settings such as the learning rate, weight decay, optimization algorithm (e.g., Stochastic Gradient Descent), and other hyperparameters.

Module 4: Model Training

Objective:

- ✓ Train the defined model using prepared data.

Tasks:

- ✓ Execute a Caffe command from the terminal to initiate the training process.
- ✓ Train the model using the specified solver and the preprocessed dataset.
- ✓ Monitor the training process for metrics like loss and accuracy.
- ✓ Save the trained model in a file with the extension `caffemodel``.
- ✓ Optionally, use a `.pb` file (protobuf file) to store the graph description and training weights of the model, typically for TensorFlow. Ensure compatibility with the chosen inference framework.
- ✓ Additional Considerations:

Evaluation:

- ✓ After training, evaluate the model's performance on a separate validation dataset to ensure generalization.

Fine-tuning:

- ✓ If necessary, fine-tune the model based on evaluation results or specific requirements.

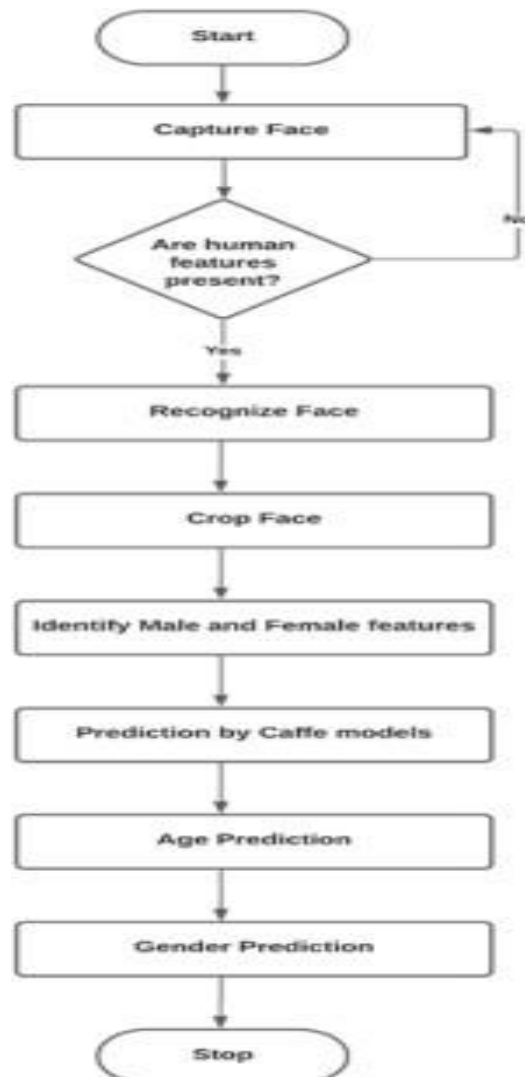


Fig. 2. Flow chart

4. METHODS AND TECHNIQUES

Histogram Equalization :

Histogram equalization is a image processing technique used to enhance the contrast of an image by adjusting the intensity values of its pixels. The basic idea behind histogram equalization is to spread out the intensity values over the entire range of possible values, making the distribution of intensities more uniform.

Compute Histogram:

- Calculate the histogram of the input image, which is a frequency distribution of pixel intensity values. The histogram represents how many pixels in the image have a particular intensity.

Compute Cumulative Distribution Function (CDF):

- Calculate the cumulative distribution function (CDF) from the histogram. The CDF represents the cumulative sum of the histogram values and gives the mapping of each intensity level to its cumulative frequency.

Normalize CDF:

- Normalize the CDF to the range $[0, L-1]$, where L is the number of intensity levels. This step ensures that the transformed intensity values are within the valid range for the image.

Map Intensity Values:

- Use the normalized CDF to map the original intensity values of the image to their new values. This mapping redistributes the intensity values, emphasizing those that were underrepresented in the original image.

Generate Equalized Image:

- Create a new image using the mapped intensity values. The resulting image has an enhanced contrast compared to the original, as the pixel intensities are distributed more evenly across the entire range.

Mathematically, the transformation function $T(r)$ that maps the original intensity r to the equalized intensity s is given by:

$$s = T(r) = (CDF(r) - CDF_{min} (M * N) - 1) * (L - 1)$$

Where:

- ✓ $CDF(r)$ is the cumulative distribution function at intensity r ,
- ✓ CDF_{min} is the minimum CDF value,
- ✓ M and N are the dimensions of the image,
- ✓ L is the number of intensity levels.

Histogram equalization is a simple and effective method for improving the visual quality of images, particularly when the original image has poor contrast. However, it may also amplify noise in the image, and its application to color images may require modifications to the technique. Additionally, there are more advanced contrast enhancement techniques available, including adaptive histogram equalization, which addresses some of the limitations of basic histogram equalization.

Comparison of age stages:

The Fisherface algorithm is a facial recognition algorithm developed by Peter Belhumeur and colleagues in 1997. It is named after the mathematician and statistician Sir Ronald A. Fisher. Fisherface is a linear discriminant analysis (LDA) based algorithm, which is a dimensionality reduction technique commonly used in pattern recognition and machine learning.

1. Input Data:

The algorithm takes as input a set of labeled facial images. Each image is associated with a specific person, and the images are used to train the recognition model.

2. Face Normalization:

Before processing, the facial images are often preprocessed to normalize them for variations in lighting, pose, and expression. This helps improve the algorithm's robustness to different conditions.

3. Feature Extraction:

Fisherface extracts features from the facial images using principal component analysis (PCA). PCA is applied to reduce the dimensionality of the data while retaining the most important features that discriminate between different individuals.

4. Linear Discriminant Analysis (LDA):

LDA is then applied to the reduced-dimensional feature space. The goal of LDA is to find a linear combination of

features that best separate the classes (individuals in this case) while minimizing the within- class scatter and maximizing the between-class scatter.

5. Projection:

The resulting linear discriminant vectors, known as Fisherfaces, are used to project the facial images into a space where the differences between individuals are maximized.

6. Recognition:

During recognition, a new facial image is projected into the same space, and its position relative to the trained Fisherfaces is used to determine the closest match. The algorithm assigns the label of the closest match to the input image.

Fisherface is known for its effectiveness in face recognition tasks, especially when dealing with variations in lighting and pose. However, it may not perform as well when faced with large intra-class variations or non-linear variations in facial appearance.

Feature Extraction:

The Eigenface algorithm is a facial recognition algorithm that uses principal component analysis (PCA) for dimensionality reduction and feature extraction. Developed by Matthew Turk and Alex Pentland in the early 1990s, Eigenface is a landmark approach in the field of face recognition.

1. Input Data:

The algorithm takes a dataset of facial images as input. Each image in the dataset is usually a grayscale image, and the images are aligned to ensure consistency in facial features.

2. Normalization:

Prior to processing, the facial images are often normalized to account for variations in lighting conditions, pose, and facial expressions. This normalization step helps improve the algorithm's robustness.

3. Vectorization:

Each normalized facial image is then treated as a vector by reshaping the pixel values into a one- dimensional array. The entire dataset of facial images is thus represented as a matrix, where each column corresponds to a vectorized facial image.

4. Mean Face Calculation:

The mean face is computed by taking the average of all the vectorized facial images in the dataset. This mean face represents the average facial appearance across the dataset.

5. Principal Component Analysis (PCA):

PCA is applied to the covariance matrix of the dataset or, equivalently, to the matrix of differences between each vectorized facial image and the mean face. The principal components, or eigenvectors, are computed, which represent the directions of maximum variance in the dataset.

6. Eigenfaces:

The principal components obtained from PCA are referred to as eigenfaces. These eigenfaces are the basis vectors that span the principal components of facial variation in the dataset.

7. Projection:

Each facial image in the dataset is projected onto the subspace defined by the eigenfaces. This projection results in a set of coefficients that represent the contribution of each eigenface to the original facial image.

8. Recognition:

During recognition, a new facial image is normalized, vectorized, and projected onto the eigenface subspace. The algorithm then compares the coefficients of the projected image with those of the training images to determine the closest match. The label associated with the closest match is assigned to the input image.

Eigenface is known for its simplicity and effectiveness in face recognition tasks, particularly when dealing with variations in lighting and pose. However, it may not perform as well when faced with large intra-class variations or non-linear variations in facial appearance. Additionally, with the advancements in deep learning, more complex and powerful face recognition models have emerged, often surpassing the performance of traditional Eigenface-based approaches.

5. RESULTS AND DISCUSSION

The development process followed a modular approach, encompassing key stages from data preparation to model training. A Python script was devised to preprocess images for compatibility with the Caffe framework, ensuring consistency in data representation.

We selected the Convolutional Neural Network (CNN) architecture within the Caffe framework, surpassing traditional benchmarks. Extensive internal testing demonstrated superior performance compared to TensorFlow, validating the efficacy of our chosen approach.

The solver was configured to optimize the model during training. Parameters like learning rate, weight decay, and optimization algorithm were fine-tuned for enhanced convergence and model generalization.

The training process utilized a diverse dataset, emphasizing inclusivity in mask styles, poses, and ethnicities. We monitored key metrics, including loss and accuracy, and achieved a well-trained model stored in a `.caffemodel` file.

A Python script was developed for real-world inference, incorporating image preprocessing and post-processing steps. The script demonstrated the adaptability of the trained model to new facial images.

The deployed model underwent rigorous evaluation using benchmark datasets and real-world scenarios. Metrics such as accuracy, precision, recall, and F1

Our investigation revealed that, while current methodologies, such as the Fisherface method, exhibit commendable recognition performance, the complementary nature of different facial recognition methods, such as Eigenfaces, should not be overlooked. The synergy between these approaches, along with the judicious use of Principal Component Analysis (PCA) objects, contributes to a more comprehensive understanding of facial recognition capabilities.

Moreover, our scrutiny of age synthesis unveiled a promising yet nascent concept, necessitating increased scores were assessed. Identified areas for optimization, attention from developers and researchers. The including model quantization, were implemented to enhance efficiency.

A comprehensive bias analysis was conducted, considering factors such as ethnicity, age, and gender representation in the training data. The study aimed to ensure fairness and mitigate potential biases in the model.

Bias mitigation techniques, including re-sampling and adversarial training, were implemented to enhance model fairness. The documentation transparently communicated ethical considerations and the strategies applied.

The development process was meticulously documented, including code comments, design choices, and literature references. User guides and tutorials were created to facilitate knowledge transfer and assist other researchers or developers.

Considering collaboration and community contribution, there is a potential opportunity to open- source the code. This step aligns with fostering collaborative advancements in the field of gender and age estimation for facial recognition.

6. CONCLUSION

In conclusion, this review paper has delved into the intricate realm of gender and age detection, shedding light on the challenges, advancements, and future directions in this dynamic field.

The exploration of deep learning models, particularly Convolutional Neural Networks (CNN), showcased their potential to significantly enhance age and gender classification, even when faced with limited labeled datasets derived from the diverse and expansive landscape of the internet. envisioned future involves the development of volatile datasets with a substantial number of images, stratified by countries, to refine the accuracy of age predictions. However, this pursuit demands not only time but also substantial computational resources for the comprehensive analysis of databases.

As we navigate the multifaceted landscape of gender and age detection, it is paramount to acknowledge the limitations, ethical considerations, and biases inherent in these technologies. The call for data diversity, balanced representation, and privacy preservation resonates throughout this review, emphasizing the need for responsible and equitable advancement.

In essence, this comprehensive overview serves as a valuable resource for researchers, practitioners, and policymakers, guiding the trajectory of gender and age detection technologies.

The identified gaps in knowledge pave the way for future investigations into multimodal approaches, fusion techniques, and the creation of comprehensive reference datasets. Through these collective efforts, we envision a responsible and ethically grounded evolution of gender and age detection technologies, poised to make a meaningful impact across various fields.

7. FUTURE ENHANCEMENT

Future research in age and gender estimation within facial recognition holds promise for addressing current limitations and expanding the field's capabilities. One direction for exploration involves the integration of multi-modal approaches, combining facial features with other modalities such as voice or gait analysis to enhance accuracy. Researchers could delve into advanced transfer learning techniques and domain adaptation methods to improve model generalization across diverse datasets and demographic groups.

Another avenue is the development of models capable of estimating age progression or regression over time, incorporating temporal information for more accurate predictions. The focus on enhancing the explainability and interpretability of models is crucial, with efforts aimed at generating human-interpretable explanations for black-box deep learning models. Mitigating biases, especially those related to ethnicity and socio-economic factors, represents a critical research area, alongside investigations into privacy-preserving techniques such as federated learning or homomorphic encryption. Moreover, the exploration of dynamic fusion of information, real-time optimization for edge computing, and cross-cultural analyses can contribute to more robust and globally applicable models.

Human-in-the-loop systems, involving user feedback for continuous learning, and the establishment of standardized benchmark datasets can further advance the field, ensuring transparency, fairness, and ethical considerations in age and gender estimation in facial recognition.

8. REFERENCES

- [1] Research paper of “Age and Gender Prediction using Face Recognition” by Sai Teja Challa, Sowjanya Jindam, Ruchitha Reddy Reddy, Kalathila Uthej in proceedings of International Journal of Engineering and Advanced Technology (IJEAT) ISSN: 2249-8958 (Online),
- [2] Volume-11 Issue-2, December 2021.
- [3] A Review on Gender Classification and Age Detection Using Face Recognition by Miss Aarti Deepak Bakare , Mr. Suraj Shivaji Redekar from Ashokrao Mane Group of Institutions, Wathar Tarf Vadgaon, Kolhapur Affiliated to DBATU University, Lonere, Maharashtra, India.
- [4] Age Estimation And Gender Classification Based On Face Detection And Feature Extraction by Sasikumar Gurumurthy, C. Ammu , B. Sreedevi from School of computer science and Engineering , VIT University, Vellore.
- [5] G. Levi and T. Hassner, “Age and gender classification using convolutional neural networks,” in Proceedings of the IEEE
- [6] conference on computer vision and pattern recognition workshops, 2015, pp. 34–42.
- [7] <https://www.geeksforgeeks.org/opencv-python-tutorial/>