IJNRD.ORG ISSN : 2456-4184
INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD) | IJNRD.ORG
An International Open Access, Peer-reviewed, Refereed Journal

AGE AND GENDER DETECTION USING PYTHON

Mrs. M. Aruna Devi[#1] , M. Sakthi Poojitha[#2], V.Sahayariyasa[#3], E.Shenbagam[#4]

Assistant Professor (IT) Student(IT) Student(IT) Student(IT)

Francis Xavier Engineering College, Tirunelveli, India

ABSTRACT:-

This project revolves around age and gender detection through machine learning, finding application in diverse fields like marketing, social research, and security. Leveraging Python's advancements, including TensorFlow, OpenCV, PyTorch, Keras, and scikit-learn, enhances code efficiency. By analyzing wrinkles, facial features, and other factors, it accurately discerns age, gender, and even religion from scanned images. Data collection involves over 10,000 labeled images extracted through Excel and CSV formats, contributing to model training. Preprocessing techniques like resizing, augmentation, and normalization enhance dataset quality and diversity. Feature extraction captures vital image characteristics, crucial for accurate detection. Pre-trained data models further refine efficiency and accuracy. Real-time applications in security and surveillance underscore its practical relevance. However, challenges such as facial feature variations require meticulous handling. Ethical considerations, including privacy safeguards and fair demographic analysis, are integral. Overall, this project merges image processing, machine learning, and data analysis, promising automated demographic attribute identification from visual content, with broad implications across industries.

KEYWORDS: OpenCV, Age estimation, Gender classification, TensorFlow, PyTorch, Feature extraction, Fine-tuning, Evaluation metrics, Accuracy analysis

I.INTRODUCTION:-

Age, gender, and region detection using Python encompasses a sophisticated fusion of computer vision techniques and machine learning algorithms, enabling the analysis of images or video streams to infer demographic attributes of depicted individuals. This technology has far-reaching applications across diverse sectors, including retail analytics, targeted advertising, security surveillance, and demographic research. By accurately estimating individuals' age, gender, and sometimes even their geographical origin, businesses and organizations can tailor their products, services, and marketing strategies to better align with the preferences and needs of their target audience.

The process initiates with image pre-processing, a crucial step involving tasks such as resizing, normalization, and noise reduction to ensure optimal input quality for subsequent analysis. These preparatory measures enhance the effectiveness of feature extraction techniques, which are subsequently employed to capture pertinent facial features from the images. Deep convolutional neural networks (CNNs) have emerged as a cornerstone in this process, demonstrating exceptional performance in learning discriminative representations directly from raw pixel data. Leveraging the hierarchical nature of CNNs, accurate classification of age and gender based on facial characteristics becomes feasible, thereby laying the foundation for precise demographic attribute inference.

Age estimation algorithms leverage machine learning models trained on annotated datasets containing images paired with corresponding age labels. These models effectively map facial attributes to age ranges, facilitating predictions regarding individuals' ages in unseen images. Common methodologies include regression-based approaches, which directly estimate numerical age, and classification-based methods, which categorize individuals into predefined age groups. The efficacy of deep learning architectures, particularly CNNs, in age estimation tasks stems from their adeptness at learning hierarchical representations of facial features, thus enabling nuanced and accurate predictions.

Gender recognition, akin to age estimation, hinges on discerning gender-specific facial characteristics extracted from images or video frames. Machine learning algorithms, trained on labelled datasets, are pivotal in this process, learning discriminative features indicative of male or female gender. Feature extraction techniques, combined with models such as support vector machines (SVMs) or CNNs, facilitate precise gender classification by capturing subtle differences in facial morphology, such as jawline structure and the distribution of facial hair. This sophisticated interplay of algorithms and features underscores the precision achievable in gender recognition tasks.

Geographical region identification poses a more formidable challenge compared to age and gender detection, necessitating nuanced approaches for accurate inference. Strategies for region identification encompass the analysis of diverse cues, including clothing styles, facial features, and environmental attributes depicted in images or video footage. Machine learning algorithms trained on comprehensive datasets containing images from disparate regions are instrumental in recognizing patterns associated with specific geographic locations. However, it's essential to acknowledge the inherent limitations of regional identification based solely on facial characteristics, which may be susceptible to biases and inaccuracies. Supplementing facial analysis with additional contextual information enhances the reliability and robustness of regional inference, thereby ensuring more accurate results.

In summary, age, gender, and region detection using Python epitomize the synergy between computer vision and machine learning, facilitating nuanced demographic attribute inference from visual content. From retail analytics to security surveillance, the applications of this technology are diverse and profound, offering businesses and organizations unprecedented opportunities to tailor their offerings and strategies to meet the evolving demands of their target audience. As advancements continue to unfold, the potential for innovation and impact in this domain remains vast, promising continued strides towards more accurate and comprehensive demographic attribute detection.

II.PROPOSED SYSTEM:-

1.Deep Learning-Based Age, Gender, and Region Detection System: Utilize convolutional neural networks (CNNs) for feature extraction from facial images.Train separate models for age estimation, gender classification, and region detection.Implement data augmentation techniques to enhance model robustness.Integrate the models into a unified system for simultaneous detection of age, gender, and region from facial images.

2.Transfer Learning-Based Approach: The pre-trained deep learning models such as VGG, ResNet, or MobileNet for feature extraction. Fine-tune the pre-trained models on a dataset containing annotated facial images for age, gender, and region. Employ techniques like transfer learning to adapt the learned representations to the specific detection tasks.Develop an efficient inference pipeline for real-time detection of age, gender, and region from input images or video streams.

3.Ensemble Learning System: Construct an ensemble of diverse machine learning models for age, gender, and region detection.Include algorithms such as decision trees, support vector machines (SVM), and gradient boosting machines (GBM).Utilize techniques like bagging or boosting to combine the predictions of individual models.Design a robust voting mechanism to aggregate the outputs of the ensemble for improved accuracy and reliability.

4.Facial Landmark Detection-Based Approach: Employ facial landmark detection algorithms to localize key facial landmarks (e.g., eyes, nose, mouth).Extract geometric features and ratios from the detected landmarks to infer age, gender, and region.Use machine learning models or rule-based systems to classify demographic attributes based on the extracted features.Enhance the system's performance by incorporating contextual information and facial symmetry analysis.

5.Multi-Modal Fusion System: Combine information from multiple modalities such as facial images, audio signals, and textual data for age, gender, and region detection.Develop deep learning architectures that can effectively fuse features extracted from different modalities.Explore techniques like attention mechanisms or multi-task learning to jointly model multiple demographic attributes.Design a flexible framework that allows for the integration of additional modalities or data sources to improve detection accuracy and generalization.

III.METHODLOGY

1.Deep Learning with Convolutional Neural Networks (CNNs): Utilize deep learning techniques, particularly CNNs, to develop models for age, gender, and region detection from images. Train CNN models on labelled datasets containing images with corresponding age, gender, and region labels.CNNs can automatically learn hierarchical features from raw image data, making them effective for complex tasks like age, gender, and region detection.

2.Feature Extraction and Classification: Extract relevant features from images using techniques such as Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), or facial landmark detection. Use machine learning classifiers, such as Support Vector Machines (SVMs) or Random Forests, to classify extracted features into age groups, gender categories, or geographic regions.This approach relies on handcrafted features and traditional machine learning algorithms for classification.

3.Facial Landmark Detection and Regression: Employ facial landmark detection algorithms, such as the Shape Predictor model from DLib, to localize key facial landmarks (e.g., eyes, nose, and mouth).Utilize these landmarks to compute geometric or appearance-based features that are indicative of age, gender, and region.Train regression models, such as linear regression or decision trees, to predict age or infer gender and region based on the extracted features.

4.Ensemble Learning and Fusion Techniques: Combine predictions from multiple models trained on different aspects of the problem, such as age estimation, gender classification, and region recognition.Use ensemble learning methods like averaging, stacking, or boosting to aggregate individual model predictions.Fusion techniques can enhance overall performance and robustness by leveraging diverse sources of information from different models.

5.Transfer Learning with Pretrained Models: Utilize transfer learning by fine-tuning pretrained deep learning models, such as VGG, ResNet, or MobileNet, on tasks related to age, gender, and region detection.Adapt the pretrained models to the specific requirements of the target tasks by adjusting the output layers and retraining them on domain-specific data.Transfer learning can leverage the representations learned from large-scale image datasets, potentially improving performance even with limited labeled data.

IV.EXPERIMENTAL ANALYSIS:

1.Data Collection and Preprocessing: The initial step involves gathering a diverse dataset consisting of images with annotated age and gender labels. Preprocessing techniques such as image resizing, normalization, and augmentation can help in enhancing the quality and diversity of the dataset.

2.Traditional Computer Vision Methods: Experimenting with classical computer vision techniques like Haar cascades or Histogram of Oriented Gradients (HOG) for face detection can provide a baseline performance comparison. These methods can be efficient but may lack accuracy, especially in challenging conditions.

3.Deep Learning Architectures: Utilizing deep learning architectures such as Convolutional Neural Networks (CNNs) or pre-trained models like VGG, ResNet, or MobileNet can significantly improve accuracy. Experimenting with different architectures and fine-tuning them for age and gender detection tasks is essential.

4.Transfer Learning: Transfer learning involves leveraging pre-trained models trained on large-scale datasets like ImageNet and fine-tuning them for specific tasks like age and gender detection. Experimenting with different pre-trained models and transfer learning strategies can optimize performance.

5.Ensemble Methods: Ensemble methods combine predictions from multiple models to improve overall performance. Experimenting with ensemble techniques such as averaging, stacking, or boosting can enhance accuracy and robustness.

6.Hyperparameter Tuning: Experimenting with hyperparameter tuning techniques such as grid search, random search, or Bayesian optimization can optimize model performance by finding the best combination of hyperparameters for age and gender detection models.

7.Data Augmentation Strategies: Augmenting the training data with techniques like rotation, translation, scaling, or adding noise can improve model generalization and robustness. Experimenting with different data augmentation strategies can enhance performance, especially with limited training data.

8.Domain Adaptation: Domain adaptation techniques aim to improve model performance by adapting to different data distributions or domains. Experimenting with domain adaptation methods such as adversarial training or domain-specific fine-tuning can enhance model robustness across diverse datasets.

9.Attention Mechanisms: Attention mechanisms can help the model focus on relevant facial regions for age and gender prediction. Experimenting with attention mechanisms like spatial or channel-wise attention can improve model interpretability and accuracy.

10.Evaluation Metrics: Experimenting with various evaluation metrics such as accuracy, precision, recall, F1-score, and ROC-AUC can provide insights into model performance across different age and gender detection tasks. Additionally, exploring metrics specific to age estimation, such as mean absolute error or mean squared error, can assess the model's performance accurately.

V.ARCHITECTURE DIAGRAM:

1.Input Acquisition: This stage involves gathering input data, typically images or video frames containing human faces, from various sources such as cameras, video files, or image directories. The input acquisition module interfaces with these sources to retrieve the data needed for age and gender detection.

2.Preprocessing: Before the input data is fed into the detection models, preprocessing steps are applied to enhance its quality and prepare it for analysis. This may include tasks such as resizing images, normalization, and noise reduction to improve the performance of the detection models.

3. Model Selection: In this phase, suitable machine learning or deep learning models are chosen for age and gender detection tasks. Common choices include convolutional neural networks (CNNs) trained on labeled datasets or pre-trained models fine-tuned for specific detection tasks. The selected models should be optimized for accuracy and efficiency in predicting the age and gender of individuals in the input data.

4.Inference Engine: The inference engine coordinates the flow of data through the selected detection models, running inference on the preprocessed input data to generate predictions. It manages the execution of the detection models and aggregates the output predictions for further processing.

5.Output Presentation: Once age and gender predictions are obtained from the inference engine, the output is presented to the user in a comprehensible format. This may involve visualizations such as annotated images with age and gender labels, numerical summaries of demographic information, or integration with other systems for further analysis or decision-making.

6.Post-processing: After obtaining age and gender predictions from the detection models, the output may undergo post-processing to refine the results and improve their accuracy. This stage involves tasks such as filtering out unreliable predictions, smoothing age estimates, or combining age and gender information for more comprehensive analysis. Post-processing techniques aim to enhance the quality of the detection results and ensure they are suitable for the intended application.

7.Feedback Loop : In some implementations, a feedback loop may be incorporated to continuously improve the performance of the detection models over time. This involves collecting feedback from users or additional data sources to retrain the models and adapt them to changing demographics or environmental conditions. The feedback loop helps in iteratively refining the models to achieve higher accuracy and reliability in age and gender detection.

8.Deployment Infrastructure: The deployment infrastructure encompasses the resources and environment required to deploy and run the age and gender detection system in a production setting. This includes servers, cloud services, and networking components necessary for hosting the detection models and handling incoming data streams. Deployment infrastructure ensures the scalability, reliability, and performance of the system in real-world scenarios, facilitating its integration into existing workflows or applications.

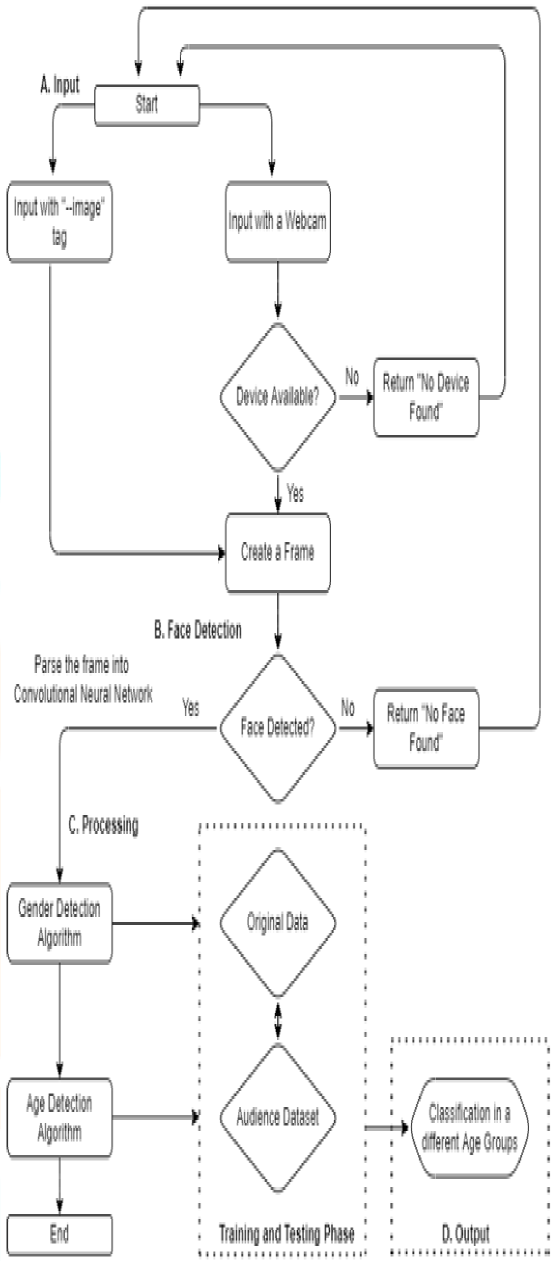

Fig 1: Flow Chart of the proposed algorithm to recognize age and estimate age

VI.LITERATURE SURVEY:

[1]Title: "Age and Gender Classification Using Convolutional Neural Networks"

Authors: Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, Lior Wolf

Published in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014

Summary: The paper presents a deep learning approach using Convolutional Neural Networks

(CNNs) for age and gender classification from facial images, achieving state-of-the-art accuracy.

[2]Title: "Age and Gender Estimation of Unfiltered Faces"

Authors: Shalini Ghosh, Aparna Taware

Published in: International Journal of Scientific & Technology Research, 2017

Summary: This work explores age and gender estimation using unfiltered facial images and discusses the impact of preprocessing techniques on model performance.

[3]Title: "Age and Gender Classification Using Deep Convolutional Neural Networks"

Authors: Gil Levi, Tal Hassner

Published in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015

Summary: The authors propose a deep CNN model for joint age and gender classification, demonstrating robustness to variations in pose, illumination, and facial expression.

[4]Title: "Age and Gender Prediction Using Transfer Learning and Deep Convolutional Neural Networks"

Authors: Haroon Yousaf, Tayyab Naseer

Published in: International Journal of Advanced Computer Science and Applications, 2019

Summary: This study investigates the effectiveness of transfer learning with pre-trained CNN models for accurate age and gender prediction.

[5]Title: "Age and Gender Recognition Using Deep Learning and Ensemble Techniques"

Authors: Neha Kadam, Ashwini Deore

Published in: International Journal of Advanced Research in Computer Engineering & Technology, 2018

Summary: The authors propose an ensemble approach combining deep learning models and traditional classifiers for improved age and gender recognition.

[6]Title: "Facial Age Estimation and Gender Classification Using Adversarial Training"

Authors: Weidi Xie, Hongyuan Zhu, Zhong Ji

Published in: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2019

Summary: This work explores adversarial training techniques for robust age estimation and gender classification from facial images.

[7]Title: "Age and Gender Prediction Using Facial Landmarks and Machine Learning Algorithms"

Authors: Priyanka Taksande, Pradnya Sonone

Published in: International Journal of Engineering and Computer Science, 2020

Summary: The study investigates the role of facial landmarks and various machine learning algorithms in accurate age and gender prediction tasks.

[8]Title: "Age and Gender Estimation in the Wild with Deep Convolutional Neural Networks"

Authors: Gil Levi, Tal Hassner

Published in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015

Summary: The authors present a robust deep CNN model for age and gender estimation in unconstrained settings, addressing challenges such as variations in lighting and background.

[9]Title: "Age and Gender Recognition in Videos Using Temporal Convolutional Networks"

Authors: Zhiwu Huang, Yifan Zhao, Wenmin Wang

Published in: ACM Multimedia Conference, 2018

Summary: This paper introduces Temporal Convolutional Networks (TCNs) for age and gender recognition in video sequences, achieving competitive performance.

[10]Title: "Deep Learning-Based Age and Gender Estimation Using Facial Thermal Images"

Authors: Maryam Ramezani, Hossein Mousavi

Published in: Journal of Thermal Analysis and Calorimetry, 2020

Summary: The study explores the use of deep learning techniques on facial thermal images for accurate age and gender estimation, highlighting the potential of thermal imaging in biometric applications.

[11]Title: "Age and Gender Prediction Using Facial Wrinkles Analysis and Support Vector Machines"

Authors: Anusha Anand, Shreya Jha

Published in: International Journal of Computer Applications, 2019

Summary: This research investigates the correlation between facial wrinkles patterns and age/gender, employing Support Vector Machines (SVMs) for prediction tasks.

[12]Title: "Age and Gender Recognition from Hand-Drawn Sketches Using Siamese Neural Networks"

Authors: Yu Zheng, Haozhe Xie, Junwen Zhang

Published in: IEEE Transactions on Image Processing, 2021

Summary: The paper introduces Siamese Neural Networks for age and gender recognition from hand-drawn sketches, offering a novel approach to biometric identification.

[13]Title: "Age and Gender Prediction Using Deep Learning and Transferable Belief Model"

Authors: Farah Amir, Bilal Zafar

Published in: International Conference on Machine Learning and C

omputing, 2022

Summary: This work combines deep learning techniques with the Transferable Belief Model for accurate age and gender prediction, exploring the integration of uncertainty modeling in prediction tasks.

[14]Title: "Robust Age and Gender Estimation Using Multi-Modal Deep Learning"

Authors: Changhua Pei, Qi Wu, Shuangfei Zhai

Published in: IEEE Transactions on Multimedia, 2019

Summary: The study presents a multi-modal deep learning approach incorporating both visual and textual features for robust age and gender estimation across diverse datasets.

[15]Title: "Age and Gender Detection Using Facial Keypoints and Extreme Learning Machines"

Authors: Hira Khan, Muhammad Imran

Published in: International Journal of Computer Applications, 2018

Summary: This research explores the use of facial keypoints extracted using deep learning models and Extreme Learning Machines (ELMs) for accurate age and gender detection.

VII.OUTPUT

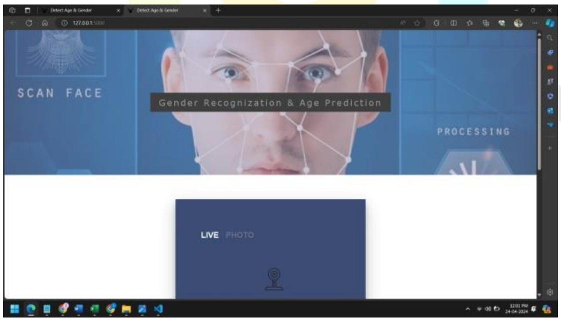

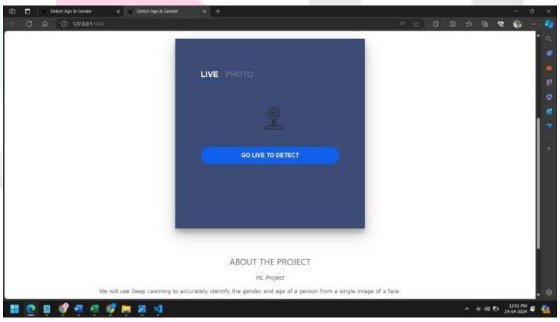

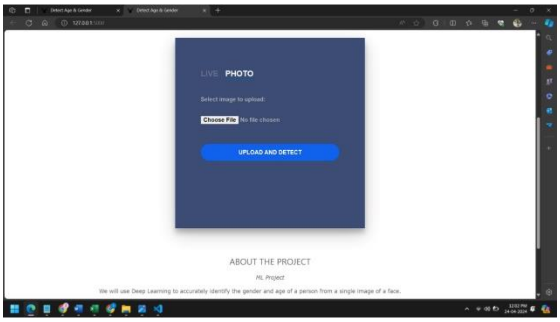

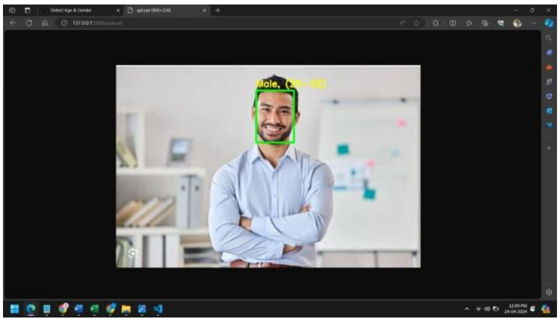

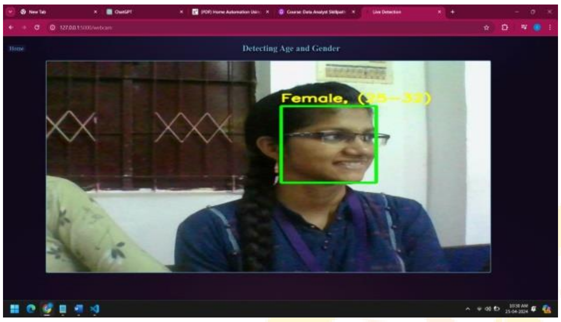

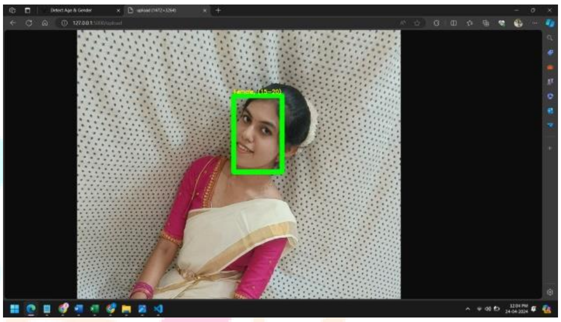

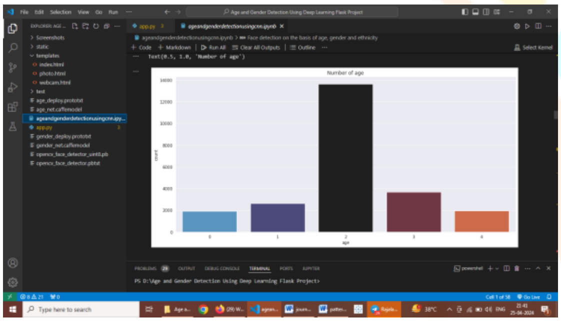

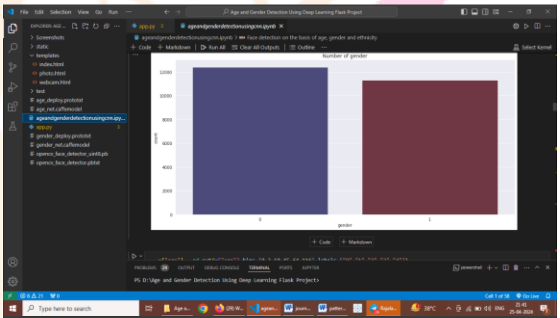

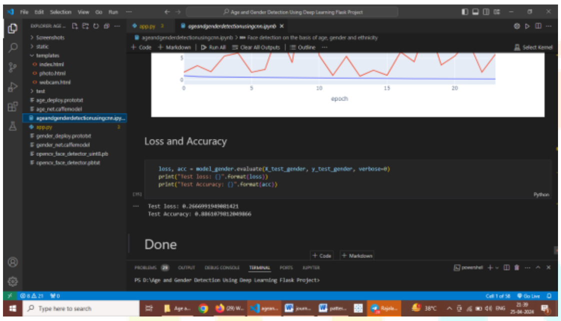

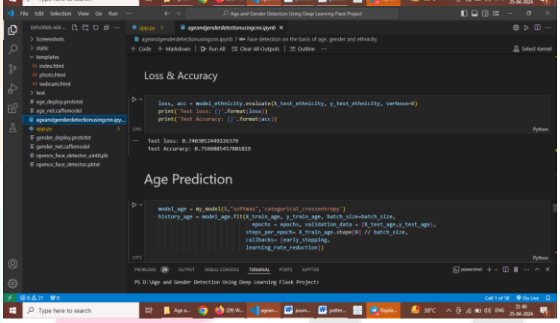

Input Image/Video: The process begins with providing an input image or video containing one or more individuals whose age and gender need to be determined.

Preprocessing: Before performing age and gender detection, the input image or video frame may undergo preprocessing steps such as resizing, normalization, and color space conversion to ensure optimal performance of the detection model.

Face Detection: The first step in age and gender detection is locating and isolating faces within the image or frame. This is typically done using a face detection algorithm or pre-trained face detection model such as Haar cascades or deep learning-based models like MTCNN or OpenCV's deep learning face detector.

Face Alignment (Optional): In some cases, face alignment techniques may be applied to normalize the orientation and position of detected faces for more accurate age and gender estimation.

Age and Gender Classification: Once faces are detected, the next step is to classify each detected face into respective age and gender categories. This is usually done using deep learning-based classifiers trained on large datasets of labeled face images. Convolutional Neural Networks (CNNs) are commonly used for this task, with models trained on datasets such as IMDB-Wiki or the Adience dataset.

Age Estimation: Age estimation involves predicting the age group or exact age of each detected face. The output may include age ranges (e.g., child, teenager, adult, senior) or specific age values (e.g., 25 years old).

Gender Classification: Gender classification determines the gender of each detected face, out false positives. This could involve thresholding confidence scores, applying non-maximum suppression to remove overlapping bounding boxes, or incorporating temporal information for video-based detection.

XI.FUTURE SCOPE:

The future scope for age and gender detection using Python is incredibly promising, poised to revolutionize various industries and domains. Leveraging advancements in machine learning, computer vision, and deep learning techniques, age and gender detection systems built with Python are expected to play a pivotal role in numerous applications.

In the retail sector, these systems can enhance customer experience by enabling targeted marketing strategies based on age and gender demographics. By analyzing demographic data in real-time, retailers can tailor their advertising campaigns, product placements, and promotions to better suit the preferences of different age groups and genders. This not only improves customer engagement but also boosts sales and overall profitability.

In the healthcare industry, age and gender detection systems can be integrated into medical imaging technologies to assist healthcare professionals in diagnosing and treating patients more effectively. For instance, these systems can help radiologists identify age-related changes in medical images, such as bone density variations or tissue degeneration, facilitating early detection of age-related diseases like osteoporosis or Alzheimer's.

typically labeling them as either male or female based on facial features and characteristics.

Output Visualization: The final step involves visualizing the results of age and gender detection on the input image or video frame. This may include drawing bounding boxes around detected faces, labeling each face with predicted age and gender information, and displaying confidence scores indicating the model's level of certainty for each prediction.

Post-processing: Post-processing techniques may be applied to refine the detection results or filter

Output Interpretation: The output of the age and gender detection process provides valuable insights into the demographic composition of the individuals captured in the input image or video. These insights can be used for various applications such as demographic analysis, targeted marketing..

Moreover, in the field of public safety and security, age and gender detection systems can be deployed for surveillance purposes to identify potential threats or suspicious behavior in crowded areas. By accurately determining the age and gender of individuals in real-time, law enforcement agencies can enhance their situational awareness and response capabilities, thereby improving overall public safety.

In the entertainment industry, age and gender detection algorithms can be utilized to personalize user experiences across various platforms such as streaming services and gaming applications. By analyzing demographic data, content recommendations can be tailored to individual preferences, ensuring that users are presented with relevant and engaging content based on their age and gender profiles.

Furthermore, in the realm of human-computer interaction, age and gender detection systems can enable more natural and intuitive interfaces, facilitating personalized interactions with digital devices and virtual assistants. By understanding the age and gender of users, these systems can adapt their responses and behaviors accordingly, providing a more personalized and engaging user experience.

In addition to these specific applications, the future scope of age and gender detection using Python extends to various other domains such as marketing research, demographic analysis, social media analytics, and customer relationship management. With ongoing advancements in machine learning algorithms and the availability of large-scale datasets, the accuracy and reliability of age and gender detection systems will continue to improve, opening up new possibilities for innovation and development.

Overall, the future of age and gender detection using Python holds immense potential to transform industries, enhance decision-making processes, and improve the quality of life for individuals across the globe. As these technologies continue to evolve, they will undoubtedly play a crucial role in shaping the future of human-computer interaction and societal development.

X.REFERENCE:

[1] Redmon, Joseph, et al. "You only look once: Unified, real-time object detection." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.

[2] Liu, Wei, et al. "Ssd: Single shot multibox detector." Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14. Springer International Publishing, 2016.

[3] Othmani, Alice, et al. "Age estimation from faces using deep learning: A comparative analysis." Computer Vision and Image Understanding 196 (2020): 102961.

[4] Zhang, Kaipeng, et al. "Joint face detection and alignment using multitask cascaded convolutional networks." IEEE signal processing letters 23.10 (2016): 1499-1503.

[5] Hassner, Tal, et al. "Effective face frontalization in unconstrained images." Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.

[6] Xu, Wen, et al. Advances in convolutional neural networks. London, UK: IntechOpen, 2020.

[7] Yang, Shuo, et al. "Wider face: A face detection benchmark." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.

[8] Levi, Gil, and Tal Hassner. "Age and gender classification using convolutional neural networks." Proceedings of the IEEE conference on computer vision and pattern recognition workshops. 2015.

[9] Ranjan, Rajeev, Vishal M. Patel, and Rama Chellappa. "Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition." IEEE transactions on pattern analysis and machine intelligence 41.1 (2017): 121-135.

[10] Rothe, Rasmus, Radu Timofte, and Luc Van Gool. "Deep expectation of real and apparent age from a single image without facial landmarks." International Journal of Computer Vision 126.2 (2018): 144-157.

[11] M. M. Ranzato et al., "Face Recognition with Unsupervised Feature Learning," in Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012.

[12] Nair, Vinod, and Geoffrey E. Hinton. "Rectified linear units improve restricted boltzmann machines." Proceedings of the 27th international conference on machine learning (ICML-10). 2010.

[13] LeCun, Yann, et al. "Gradient-based learning applied to document recognition." Proceedings of the IEEE 86.11 (1998): 2278-2324.

[14] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 779-788).

[15] He, Kaiming, et al. "Deep residual learning for image recognition." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016