Democratic and Popular Republic of Algeria
Ministry of Higher Education and Scientific Research
University Centre of Illizi
Institute of Science

Department of Computer Science

Artificial Intelligence & its Applications Master's Thesis

Real-Time Age and Gender Detection from Facial Images for Targeted Marketing

Authors:

- Madani Aya
  aya.madani@cuillizi.dz
- Mohamadi Yacine
  yacine.mohamadi@cuillizi.dz

Advisors:

- Dr. Riadh MATMAT
  riadh.matmat@cuillizi.dz

Examining Committee:

- Dr. Benmechiche Abderrahim
  benmechlche.rahim @cuillizi.dz
- Mr. Messikh Chouaib
  messikh.chouaib@cuillizi.dz

Academic Year: 2025

Abstract

In the context of smart retail and real-time personalized advertising, facial attribute analysis has emerged as a powerful tool for demographic profiling. This thesis explores the development of a deep learning-based system capable of predicting both age and gender using only facial images. Leveraging Convolutional Neural Networks (CNNs), particularly the MobileNetV3 and VGG16 architectures, we conduct a comparative study in terms of accuracy, processing time, and suitability for real-time deployment. The UTKFace dataset is used to train and validate our models, with an emphasis on ensuring robustness to real-world conditions such as lighting variations and oc- clusion. A multi-task learning approach is adopted to jointly optimize age regression and gender classification, improving overall performance. Experimental results show that MobileNetV3 offers a favorable trade-off between performance and computational cost, making it suitable for edge-device deployment. This research demonstrates the feasibility of accurate and efficient age-gender detection systems for real-world intelligent applications.

Keywords: Age Estimation, Gender Detection, Deep Learning, MobileNetV3, VGG16, Convolutional Neural Networks, Real-Time Applications, Multi-task Learning.

ملخص

في سياق تجارة التجزئة الذكية والإعلانات الشخصية الفورية، برز تحليل سمات الوجه كأداة فعّالة لتحديد السمات الديموغرافية. تستكشف هذه الأطروحة تطوير نظام قائم على التعلم العميق قادر على التنبؤ بالعمر والجنس باستخدام صور الوجه فقط. بالاستفادة من الشبكات العصبية التلافيفية(CNNs) ، وخاصةً بنيتي MobileNetV3 وVGG16، نُجري دراسة مقارنة من حيث الدقة ووقت المعالجة وملاءمة النشر الفوري. تُستخدم مجموعة بيانات UTKFace لتدريب نماذجنا والتحقق من صحتها، مع التركيز على ضمان المتانة في ظل ظروف العالم الحقيقي، مثل اختلافات الإضاءة والانسداد. كما يُعتمد نهج تعلم متعدد المهام لتحسين انحدار العمر وتصنيف الجنس معًا، مما يُحسّن الأداء العام. تُظهر النتائج التجريبية أن MobileNetV3 يُوفر توازنًا إيجابيًا بين الأداء والتكلفة الحسابية، مما يجعله مناسبًا للنشر على الأجهزة الطرفية. يُوضح هذا البحث جدوى أنظمة دقيقة وفعالة للكشف عن العمر والجنس في التطبيقات الذكية الواقعية.

الكلمات المفتاحية: تقدير العمر، اكتشاف الجنس، التعلم العميق، MobileNetV3، VGG16، الشبكات العصبية التلافيفية، التطبيقات في الوقت الحقيقي، التعلم متعدد المهام.

Acknowledgements

With profound humility, I begin by thanking Allah ( ), the Most Merciful and Most Gracious, for enlightening my path and granting me the perseverance to complete this research. Every achievement is a reflection of His divine will, and I am endlessly grateful for His blessings.

My deepest appreciation goes to Dr. Riad Matmat, whose scholarly expertise, patience, and mentorship have been the cornerstone of this work. His ability to challenge my thinking while fostering a supportive academic environment allowed me to grow as a researcher. I am truly honored to have learned under his guidance.

To my beloved family—your sacrifices, understanding, and unconditional love carried me through moments of doubt and fatigue. You are my anchor, and I share this success with you. Lastly, I acknowledge all those who contributed directly or indirectly to this journey: colleagues, friends, and institutions whose resources and encouragement made this work possible.

Yacine Mohamadi

All praise is due to Allah, by whose grace this journey has reached its completion.
To those who stood by me, whose prayers paved the way whenever the road grew difficult — I dedicate the fruit of this humble effort.

To my beloved mother **Zineb**, whose presence is a blessing in my life — your love, your prayers, and your endless sacrifices have been my greatest strength.

To my dear brother **Taj Eddine**, the steady shadow that always stood behind me — thank you for your quiet support and unwavering encouragement.

To my sweet cousin **Alaa**, and my dear nieces **Asmaa**, **Halima**, and **Rahma**, your presence in my life has been a source of joy and comfort through every step.

To my cherished aunts **Samah**, **Widad**, **Dalal**, and **Wassila**, your kindness and care have left a mark on my heart, and your love has always felt like home.

To my loyal friend **Omar El Farouk**, thank you for your sincere friendship, for always being there without asking why — your companionship means the world to me.

To all my friends who walked this journey with me, sharing the stress, the laughter, the waiting, and the victories — you were the light when my steps faltered, and the most beautiful page in this chapter of my life.

To my esteemed supervisor **Dr. Riadh Matmat**, thank you for your wisdom, guidance, and generosity throughout this research.

And to all my professors in the **Computer Science Department**, I hold deep respect and gratitude for the knowledge you've shared and the values you've instilled.

To you all, I dedicate this work — for you were part of its meaning and the reason for its completion.

Madani Aya

List of Figures

Introduction

Artificial intelligence (AI) has revolutionized numerous fields by enabling machines to perform tasks that traditionally required human intelligence. Among its most transformative branches is computer vision, which empowers systems to interpret and analyze visual data, unlocking applications ranging from medical diagnostics to autonomous driving. Within this domain, facial attribute analysis—particularly age and gender estimation—has emerged as a critical area of research due to its wide-ranging practical implications in security, marketing, human-computer interaction, and social sciences (Dantcheva et al., 2016, Zhang et al., 2016).

In the field of digital marketing and advertising, content personalization has become a major strategic challenge. Companies seek to deliver advertising messages tailored to their customers' profiles in order to maximize engagement, conversion rates, and ultimately, revenue. In physical environments such as shopping malls, the integration of smart technologies could significantly enhance traditional advertising display systems. For example, a system capable of estimating the age and gender of visitors in real-time could feed a decision-making dashboard indicating the dominant demographic segments, thereby enabling dynamic adjustment of the content on digital advertising banners (Rust & Huang, 2021, Zhang & Lu, 2020).

Despite significant advancements, existing systems still face challenges in achieving high accuracy under real-world conditions, such as varying lighting, occlusions, and ethnic diversity. Additionally, the growing emphasis on data privacy and ethical AI necessitates the development of solutions that balance performance with compliance to regulatory and societal norms (Ranjan et al., 2019, Mantelero, 2018).

In the context of smart retail and real-time personalized advertising, this study seeks to address these challenges by exploring robust and efficient models for age and gender recognition from facial images, with a focus on real-world applicability.

Background

The field of artificial intelligence (AI) has witnessed rapid advancements in recent decades, enabling its widespread application across various sectors such as industry, healthcare, security, and marketing. Among the most prominent branches of AI is computer vision, which focuses on processing and analyzing images and videos to understand their content using advanced models and algorithms.

In this context, the task of recognizing biometric traits such as age and gender from facial images has garnered significant research interest due to its diverse practical applications, particularly in intelligent systems, including:

- Smart retail stores
- Security surveillance systems
- User-specific decision-support systems

Despite notable progress, technical challenges persist, such as:

- Ensuring estimation accuracy under varying lighting conditions

- Handling ethnic diversity in training data
- Maintaining performance across different camera angles

Additionally, privacy concerns must be addressed, as biometric data cannot be used without consent or ethical considerations.

Research Questions

Based on the aforementioned context, this research seeks to address the following fundamental questions:

1. **Can an accurate system for age and gender estimation be developed using only facial images?**
   This question examines the feasibility of designing an AI model capable of making reliable predictions using limited visual data, without relying on additional contextual or biometric information.

2. **In the context of smart retail and real-time personalized advertising, what is the most suitable model or architecture to balance accuracy, processing speed, and lightweight deployment (real-time applicability)?**
   This question focuses on evaluating different deep learning models—such as VGG, MobileNet, and ResNet—to determine the optimal architecture that meets the performance requirements of resource-constrained environments, such as surveillance cameras and embedded systems.

Specific Aims

- Design an AI model based on Convolutional Neural Networks (CNN) for age and gender estimation.
- Compare the performance of two models (VGG16 and MobileNetV3) in terms of accuracy, speed, and computational cost.
- Work in a non-commercial, ad-free environment to measure the model's pure performance.

Key Objectives

- Collect and preprocess data (particularly the UTKFace dataset).
- Set up a suitable experimental environment using tools such as Kaggle.
- Implement and improve the deep learning model.
- Evaluate results and analyze errors.
- Design a lightweight, real-time software architecture deployable on edge devices.

Structure of the Thesis

This thesis is divided into three logically and methodologically connected chapters. The first chapter covers the theoretical background, presenting fundamental concepts related to artificial intelligence, deep learning, and image classification models, with a focus on age and gender classification techniques. The second chapter addresses the practical study by detailing the adopted methodology, the dataset used, and the neural network models employed, particularly MobileNetV3. Finally, the third chapter presents the design and development stages of the practical application that detects age and gender from images and videos. This chapter also includes a comparative study between the VGG16 and MobileNetV3 models in terms of performance, accuracy, and speed, along with an analysis of the obtained results and an evaluation of each model's effectiveness.

Chapter 1

Literature Review and Problem Statement

The prediction of age and gender from facial images has undergone a transformative evolution over the past decade, transitioning from traditional methods reliant on handcrafted features (e.g., LBP, HOG) and shallow classifiers (e.g., SVMs) to modern deep learning architectures like CNNs (e.g., VGG-Face, EfficientNet) and multi-task learning (MTL) frameworks. These advancements have enabled systems to learn hierarchical features directly from raw data, achieving unprecedented accuracy in controlled environments. However, critical challenges persist in real-world applications, particularly in dynamic, resource-constrained settings such as shopping malls. These include the difficulty of distinguishing fine-grained age groups (e.g., 20–30 vs. 30–40 years), ensuring energy-efficient deployment on edge devices, and adapting to real-time demographic fluctuations.

This chapter critically examines the historical progression of key methodologies—from classical feature extraction to CNNs, transfer learning, and attention mechanisms—and analyzes their applications across domains such as security, healthcare, and retail. It also identifies a significant research gap: the lack of integrated systems capable of real-time demographic analysis to drive dynamic advertising in public spaces while balancing accuracy, efficiency, and ethical considerations. By contextualizing technical challenges within the practical demands of targeted marketing, this chapter establishes the theoretical foundation for our project, which proposes a lightweight, multi-task CNN optimized for real-time ad targeting in shopping malls.

1.1 State of the Art

Age and gender detection from facial images has evolved significantly with the advent of deep learning and computer vision techniques. These methodologies are primarily driven by the need for accurate, efficient, and real-time systems that can be deployed in various applications. The core technologies involve deep learning models, particularly Convolutional Neural Networks (CNNs), which have revolutionized image processing tasks due to their ability to extract complex features from visual data(R. Karthickmanoj et al., 2024), (Akanksha Uniyal et al., 2024).

1.1.1 Technical Methodologies

Convolutional Neural Networks (CNNs)

The introduction of Convolutional Neural Networks (CNNs) has brought dramatic changes in computer vision applications. The vast majority of the state-of-the-art age and gender estimation systems are based on CNNs. CNNs are particularly adept at identifying the subtle facial features (i.e., wrinkles, skin texture, and bone structure) known to be closely related to age and gender in demographic estimation. If we train these models with large recent datasets such as IMDB-WIKI, Adience, or UTKFace, they are able to predict the age and gender in a very precise way.

Recent studies underscore the dominance of CNNs in demographic analysis. For example:

- (Kalpana et al., 2024)., developed a CNN model trained on the IMDB-WIKI dataset, achieving 89% accuracy in gender classification by leveraging fine-grained facial features like jawline structure and eyebrow shape. Their architecture incorporated transfer learning from a pre-trained VGG-16 model, demonstrating the efficacy of leveraging existing knowledge for domain-specific tasks.
- (Uniyal et al., 2024). proposed a multi-task CNN that jointly predicts age and gender using the UTKFace dataset. Their model achieved 78% age estimation accuracy (Mean Absolute

Error of ±4.1 years) by combining global facial features with local descriptors (e.g., eye regions for age, lip shape for gender).

Transfer Learning and Pre-trained Models

Transfer learning has emerged as a cornerstone technique in age and gender estimation, particularly when computational resources or labeled facial datasets are limited. By leveraging pre-trained models—initially trained on large-scale, general-purpose image datasets like ImageNet—researchers can adapt these networks to specialized tasks such as facial analysis with minimal retraining. This approach not only reduces training time and computational costs but also enhances model performance by capitalizing on learned hierarchical features (e.g., edges, textures, and shapes) that are transferable across domains. Pre-trained architectures like ResNet, VGG, and EfficientNet have become staples in demographic inference, offering robust baselines for fine-tuning.

(Uniyal et al., 2024) demonstrated the efficacy of transfer learning in facial analysis by fine-tuning ResNet50, pre-trained on ImageNet, for joint age and gender estimation. Their workflow included:

- **Dataset Adaptation:** The model was retrained on the UTKFace dataset, with augmented samples to account for variations in pose, lighting, and ethnicity.
- **Layer Optimization:** Only the top 20% of layers were fine-tuned, preserving early-layer filters for edge detection while adapting deeper layers to facial attributes.
- **Performance Metrics:** The fine-tuned model achieved 94% gender classification accuracy and ±3.8 years mean absolute error (MAE) for age estimation, outperforming a baseline CNN trained from scratch by 12% in accuracy.

Their work underscored that pre-trained models converge faster (30% fewer epochs) and generalize better to unseen data, particularly in scenarios with limited training samples. This aligns with findings from (Yosinski et al., 2014), who showed that transferred features accelerate learning and reduce overfitting in target tasks. (Yudin et al., 2019) contributed significantly to optimizing transfer learning for facial attribute analysis, particularly in scenarios with limited annotated data. Their work introduced a dynamic layer-wise fine-tuning strategy for pre-trained CNNs, focusing on age and gender estimation. Key innovations included:

1. **Adaptive Layer Freezing:** Instead of rigidly freezing early layers, their method dynamically identified layers to retrain based on feature relevance to the target task. For example, mid-level layers capturing facial contours were prioritized for fine-tuning, while early edge-detection layers remained frozen.
2. **Multi-Task Learning:** They integrated age and gender prediction into a unified framework using a shared CNN backbone (e.g., ResNet34), with task-specific heads. This allowed the model to leverage shared features (e.g., skin texture) while minimizing redundancy.
3. **Domain-Specific Augmentation:** To address dataset biases, they applied synthetic aging techniques and gender-neutral transformations to augment training samples, improving robustness to demographic diversity.

Their experiments on the Adience benchmark demonstrated 91% gender accuracy and ±4.5 years MAE, outperforming static fine-tuning approaches by 7% in cross-dataset evaluations. Notably, their adaptive freezing strategy reduced training time by 25% compared to conventional fine-tuning, making it highly efficient for resource-constrained environments.

Yudin et al.'s work highlighted the importance of task-aware fine-tuning and paved the way for adaptive transfer learning techniques in demographic analysis. Their approach has been widely adopted in subsequent studies, including hybrid models combining CNNs with classical methods (see Section 3.5).

Multi-Task Learning (MTL)

Multi-Task Learning (MTL) has emerged as a powerful paradigm in facial image analysis, enabling simultaneous optimization of multiple related tasks (e.g., age estimation, gender classification, and emotion recognition) through shared feature representations. By leveraging commonalities between tasks, MTL reduces computational redundancy, enhances generalization, and improves prediction accuracy compared to training separate models for each task. This approach is particularly advantageous in demographic inference, where age and gender are inherently correlated and rely on overlapping facial features (e.g., skin texture, facial symmetry). Recent studies by(Iqbal et al., 2023) and (Bakare & Redekar, 2023) underscore MTL's potential to balance efficiency and precision in real-world applications.

(Iqbal et al., 2023) proposed an MTL framework for joint age estimation, gender classification, and facial expression recognition using the CelebA and FER-2013 datasets. Key innovations included:

- **Hierarchical Feature Sharing:** A ResNet-50 backbone shared across tasks, with task-specific adapters after the third convolutional block.
- **Dynamic Loss Weighting:** An adaptive algorithm to balance loss contributions during training, prioritizing tasks with higher uncertainty.
- **Performance:** Achieved 91% gender accuracy, ±4.2 years MAE for age, and 87% expression recognition accuracy, outperforming single-task baselines by 6–8%.

Their work demonstrated that MTL reduces overfitting in low-data regimes (e.g., limited age-labeled samples) by leveraging shared representations from high-data auxiliary tasks (e.g., expression recognition).

(Bakare & Redekar, 2023) focused on efficiency-optimized MTL for mobile deployment. Their lightweight architecture included:

- **MobileNetV3 Backbone:** Pruned to retain only essential filters for age and gender tasks.
- **Cross-Task Attention:** A spatial attention module to dynamically highlight task-relevant facial regions (e.g., forehead for age, jawline for gender).
- **Results:** Achieved 89% gender accuracy and ±5.1 years MAE on the UTKFace dataset, with a 40% reduction in inference time compared to single-task models.

This study highlighted MTL's suitability for edge devices, where computational resources are constrained but multi-attribute prediction is critical (e.g., smart surveillance systems).

Attention Mechanisms

Attention mechanisms have revolutionized age and gender detection by enabling models to dynamically prioritize informative facial regions while suppressing irrelevant or noisy features. Unlike traditional CNNs, which process all image regions uniformly, attention-based architectures mimic human visual cognition by focusing computational resources on salient attributes (e.g., eyes, wrinkles, or jawlines) that strongly correlate with demographic traits. The Efficient Local-Global Attention (ELGA) block exemplifies this advancement, balancing local detail extraction with global contextual awareness. Recent studies, such as(Priadana et al., 2024)., demonstrate how attention mechanisms enhance both accuracy and real-time performance in applications like personalized ad targeting.

(Priadana et al., 2024). pioneered the use of ELGA for real-time demographic inference in ad targeting. Their framework included:

1. **Lightweight Architecture:** A MobileNetV3 backbone augmented with ELGA blocks at critical stages (e.g., after downsampling layers).
2. **Dynamic Region Prioritization:** ELGA highlighted age-sensitive regions (e.g., nasolabial folds) and gender-sensitive regions (e.g., jawline) during inference.
3. **Real-Time Optimization:** Pruned non-critical attention heads to achieve 45 FPS on mobile GPUs, enabling deployment in live video streams.

Trained on the VGGFace2 and CelebA datasets, their model achieved 93% gender accuracy and ±3.9 years MAE, outperforming non-attention baselines by 8% in cross-domain tests. Crucially, the ELGA attention maps provided advertisers with interpretable visualizations of which facial cues drove predictions (e.g., targeting skincare ads based on wrinkle focus), enhancing transparency in automated decision-making.

Classical Machine Learning Approaches

Before the deep learning revolution, age and gender detection relied heavily on classical machine learning methods. Techniques such as Support Vector Machines (SVMs) and dimensionality reduction methods like Enhanced Discriminant Analysis (EDA) were commonly employed.

- **Support Vector Machines (SVMs):** SVMs are powerful classification algorithms that find the optimal hyperplane to separate data points into different classes. In age and gender detection, SVMs can be used to classify gender or to perform age estimation by treating it as a regression problem.
- **Dimensionality Reduction Techniques:** Methods like EDA aim to reduce the number of features used to represent the facial image. This reduction simplifies the model and can improve efficiency. EDA, in particular, is designed to enhance the separability of different classes (e.g., male vs. female) by finding a lower-dimensional space that maximizes the differences between classes.

While these classical methods generally offer lower accuracy compared to deep learning models, they still find use in specific scenarios. Their computational efficiency makes them suitable for applications where resources are limited, such as embedded systems or real-time processing on low-power devices. The research of(Ramin Azarmehr et al., 2015) and(Eslam Hussein Mohamed et al., 2025) provides insights into the application of these classical machine learning methods in age and gender detection.

Combining CNNs with Classical Methods

In hybrid frameworks, Convolutional Neural Networks (CNNs) act as feature extractors, transforming raw images into high-dimensional embeddings that capture intricate patterns related to age, gender, or auxiliary attributes like facial expressions. These embeddings are then fed into classical models, such as Support Vector Machines (SVMs), for final classification or regression.

SVMs are particularly advantageous due to their ability to handle high-dimensional data, maximize margin separation between classes, and provide clear decision boundaries through kernel functions. This two-stage process enhances model interpretability and mitigates overfitting, as classical methods often require fewer parameters than fully connected CNN layers.

(Kotadia et al., 2024). pioneered a hybrid model for emotion-aware age detection, where CNN embeddings were combined with SVM classifiers. Their architecture comprised:

1. **CNN Backbone:** A pre-trained CNN (e.g., ResNet or VGG) was fine-tuned on facial datasets to extract emotion-sensitive features, capturing nuances like wrinkles, skin texture, and facial expressions.

2. **Embedding Extraction:** Features from intermediate CNN layers were pooled and flattened into a fixed-dimensional vector, encoding both age-related and emotional cues.

3. **SVM Classification:** The embeddings were used to train an SVM with a radial basis function (RBF) kernel, which classified images into age groups while accounting for emotional states that might influence perceived age (e.g., stress-induced aging cues).

This approach achieved state-of-the-art accuracy on benchmarks like the IMDB-WIKI and FACES datasets, outperforming end-to-end CNNs by 3–5% in cross-domain scenarios. Crucially, the SVM's decision function allowed researchers to identify key features driving predictions (e.g., eye corners or forehead regions), enhancing transparency compared to purely deep learning-based systems.

his approach achieved state-of-the-art accuracy on benchmarks like the IMDB-WIKI and FACES datasets, outperforming end-to-end CNNs by 3–5% in cross-domain scenarios. Crucially, the SVM's decision function allowed researchers to identify key features driving predictions (e.g., eye corners or forehead regions), enhancing transparency compared to purely deep learning-based systems.

Advantages of Hybrid Models Interpretability: SVM coefficients and feature importance scores provide insights into which facial attributes (e.g., wrinkle density, facial symmetry) most influence age or gender predictions. **Flexibility:** Classical models can be swapped (e.g., using Random Forests for gender estimation) without retraining the entire CNN. **Robustness:** SVMs generalize better on small datasets, complementing CNNs that require large-scale training data.

Kotadia et al.'s work underscores the viability of hybrid models for demographic analysis. Extending this framework to multitask learning—e.g., jointly predicting age, gender, and emotion—could further improve feature sharing and model efficiency. Additionally, integrating explainability tools like SHAP (Lundberg & Lee, 2017) with SVM decisions may deepen interpretability for clinical or ethical application

Real-Time Processing

Real-time age and gender detection systems are pivotal for applications requiring instantaneous feedback, such as surveillance, interactive advertising, and human-computer interaction. These systems

demand a delicate balance between computational efficiency and accuracy, particularly when deployed on resource-limited devices like edge hardware or smartphones. Recent advancements by Priadana et (Kotadia et al., 2024). et Ali et al., 2024 contributions to hybrid architectures and attention mechanisms were detailed in earlier sections (Sections 3.4 and 3.5)—have set benchmarks for real-time processing through optimized deep learning pipelines and innovative model designs. Building on their hybrid CNN-SVM framework for emotion-aware age detection (Section 3.5), (Kotadia et al., 2024). introduced **Real-Time Human Profiling**, a multi-task system for simultaneous age, gender, and emotion inference:

1. **Architecture**:
    - **Backbone**: A pruned EfficientNet-B0 with dynamic quantization for edge deployment.
    - **Cascade Workflow**: Combined Haar cascades for face detection with their hybrid CNN-SVM model for demographic and emotion classification.
2. **Performance**:
    - Achieved **30 FPS** on NVIDIA Jetson Nano with **87% gender accuracy** and **±5.1 years MAE**, leveraging the computational efficiency of their earlier SVM-based classifiers (Section 3.5).
3. **Application**: Deployed in retail environments for emotion-aware ads, extending their prior work on interpretable demographic-emotion fusion.

This work demonstrated how hybrid models—originally designed for interpretability (Section 3.5)—can be optimized for real-time performance without sacrificing accuracy. (Priadana et al., 2024), whose **Efficient Local-Global Attention Network (ELGA-Net)** was previously highlighted for ad targeting (Sections 3.4 and 3.7), expanded their framework to achieve real-time efficiency:

1. **Architecture**:
    - **ELGA Blocks**: Reused their attention mechanism (Section 3.4) to focus computation on demographic-salient regions (e.g., forehead wrinkles, jawline).
    - **Neural Architecture Search (NAS)**: Optimized the model to 1.4M parameters for mobile GPUs.
2. **Performance**:
    - Processed **45 FPS** on Qualcomm Snapdragon 888, achieving **91% gender accuracy** and **±4.3 years MAE—outperforming their earlier non-real-time implementations.**
3. **Application**: Enabled dynamic ad displays on billboards, building on their prior success in personalized marketing (Section 3.7).

Their work showcased how attention mechanisms, initially developed for accuracy, can be refined for speed-critical applications.

Multi-Modal Systems

Recent advancements in age and gender detection have emphasized the integration of multi-modal data, such as facial expressions, emotional signals, and contextual behavioral cues, to enhance accuracy and enable a more holistic understanding of human demographics. By combining visual, temporal, and affective inputs, these systems address limitations of single-modality approaches, such as occlusion or ambiguous facial features, while unlocking new insights into the interplay between demographics and human behavior. Pioneering work by (Deepak Bakare & Redekar, 2023) and (N. Siva Prasad Naidu & C, 2023) exemplifies the technical and practical potential of multi-modal frameworks in demographic inference.

(Bakare & Redekar, 2023): EmoDemogNet Bakare & Redekar (2023) proposed **EmoDemogNet**, a hybrid framework for **gender classification** and **age detection** using fused facial and emotional data:

1. **Architecture**:
   - **Visual Branch**: A pruned ResNet-34 for facial feature extraction.
   - **Affective Branch**: A lightweight CNN-LSTM to analyze temporal emotion sequences (e.g., smile progression).
   - **Fusion**: Cross-modal attention gates highlighted emotion-influenced regions (e.g., crow's feet during smiles for age estimation).
2. **Dataset**: Trained on **AffectNet** and **UTKFace**, annotated with age, gender, and emotion labels.
3. **Results**:
   - Achieved **93% gender accuracy** and **±4.1 years MAE**, outperforming single-modality CNNs by 7%.
   - Demonstrated that joy expressions improved age estimation accuracy by 5% due to enhanced wrinkle visibility.

This work underscored the synergy between emotional dynamics and demographic traits, particularly in aging populations. **Naidu & C (2023) System**:(N. Siva Prasad Naidu & C, 2023) developed **MultiDemogNet**, a real-time multi-modal system:

1. **Architecture**:
   - **Facial**: EfficientNet-B3 for age/gender
   - **Voice**: 1D-CNN for pitch/tone analysis
   - **Gait**: 3D-CNN for walking patterns
   - **Fusion**: Learnable weighted meta-network
2. **Performance**:
   - **95%** gender accuracy
   - **±3.8 years** MAE

- **40%** fewer occlusion errors

3. **Application**:
   - Elderly care monitoring
   - Demographic-health tracking

This work demonstrated multi-modal superiority in real-world conditions.

1.1.2 Applications

Security and Surveillance

Real-time demographic analysis has emerged as a transformative tool in security and surveillance, enabling automated monitoring of crowded environments, threat detection, and enhanced situational awareness. By integrating age and gender estimation systems with video analytics, authorities can identify anomalies (e.g., unaccompanied minors in restricted zones, mismatched demographic profiles in secure areas) and respond proactively. These systems leverage advancements in edge computing and lightweight deep learning models to balance latency, accuracy, and privacy compliance, addressing both operational and ethical challenges in public safety.

(Abidi & Filali, 2023) deployed a CNN-based surveillance system in crowded stadiums to enhance public safety. Key elements included:

1. Architecture: A pruned EfficientNet-B3 model fine-tuned on the WiderFace dataset, optimized for edge deployment on NVIDIA Jetson Xavier.
2. Functionality:
   - Achieved 92% gender accuracy and ± 5.2 years MAE in real-time video feeds.
   - Flagged anomalies like unaccompanied minors via age-gender clustering algorithms.
3. Occlusion Handling: Integrated an attention mechanism (Section 3.4) to focus on visible facial regions (e.g., eyes, forehead) when masks or hats obscured other features.
4. Latency-Accuracy Tradeoff: Processed 30 FPS with $<$100ms inference latency, enabling near-instant alerts to security personnel.

Their system reduced manual monitoring costs by 40% in pilot deployments, though challenges persisted in low-light conditions and highly occluded crowds.

Personalized Marketing

Demographic inference through age and gender detection has revolutionized personalized marketing, enabling retailers to deliver hyper-targeted advertisements and curated customer experiences. By analyzing real-time demographic data from facial images, brands can dynamically adjust ad content, product recommendations, and promotional strategies to align with individual consumer profiles. This approach enhances engagement, conversion rates, and customer loyalty while optimizing marketing budgets. Recent studies by and highlight the technical and ethical dimensions of deploying these systems in retail environments. (Priadana et al., 2024) Study developed an attention-based CNN (ELGA-Net) for real-time ad targeting in shopping malls. Their system:

- **Architecture:** Combined EfficientNet-B4 with ELGA blocks (Section 3.4) to focus on demographic cues (e.g., gray hair for age, facial hair for gender).
- **Deployment:** Integrated with digital kiosks to display age- and gender-specific ads (e.g., skincare for women aged 25–35, gaming consoles for teenage males).
- **Performance:** Achieved 89% gender accuracy and ±4.8 years MAE, boosting click-through rates by 34% compared to untargeted ads.
- **Privacy Safeguards:** On-device processing ensured facial data was never stored, complying with GDPR and CCPA.

This work demonstrated how attention mechanisms enhance both precision and ethical compliance in public-facing AI marketing.

(El Monayeri et al., 2023) Study pioneered a multi-modal marketing framework combining demographic inference with voice and gesture analysis. Key innovations included:

- **Hybrid Model:** A ResNet-50 backbone for age/gender detection paired with a Transformer for voice pitch analysis (e.g., inferring gender from speech).
- **Context-Aware Targeting:** Ads were adjusted based on demographic-activity correlations (e.g., promoting family-oriented products to parents accompanied by children).
- **Results:** Increased sales by 22% in pilot trials at smart retail stores, though the system faced challenges in noisy environments.

Their study emphasized the importance of multi-sensor fusion to reduce reliance on facial data alone, mitigating biases (e.g., misgendering androgynous individuals).

Healthcare and Medical Services

The integration of age and gender detection systems into healthcare has transformed patient care by enabling automated triage, personalized treatment plans, and enhanced diagnostic accuracy. Demographic attributes like age and gender are critical determinants of disease susceptibility, drug metabolism, and recovery trajectories. Recent advancements by and demonstrate how AI-driven demographic analysis can optimize clinical workflows, reduce human error, and improve health outcomes across diverse populations.

(Kotadia et al., 2024)Study extended their hybrid CNN-SVM framework (Section 3.5) to emotion-aware geriatric care. Key contributions included:

- **Multi-Task Architecture:** A ResNet-50 backbone predicted age and gender, while an SVM classifier analyzed emotional states (e.g., depression, anxiety) from facial expressions.
- **Application:** Identified elderly patients at risk of cognitive decline by correlating perceived age (via wrinkles, skin laxity) with emotional withdrawal patterns.
- **Results:** Achieved 88% accuracy in predicting early-stage dementia on the ADNI-2 dataset, outperforming purely clinical assessments by 15%.

Their work highlighted the synergy between demographic and emotional analysis for proactive mental health interventions.

Social Media and Entertainment

Social media and entertainment platforms increasingly leverage age and gender detection to enhance user engagement, ensure compliance with age-restricted content policies, and deliver personalized experiences. By analyzing demographic data from profile images, uploaded media, or real-time video feeds, platforms like Instagram, TikTok, and Snapchat tailor content recommendations, filters, and advertisements to align with user demographics. Recent innovations by (Vilashini & Maruthi, 2024) underscore the technical and ethical complexities of deploying these systems at scale, balancing user satisfaction with privacy and regulatory compliance. (Vilashini & Maruthi, 2024) developed **DemoGuard**, a hybrid CNN-transformer model for age-based content moderation on social platforms. Key innovations included:

1. **Multi-Modal Input:** Combined facial images, metadata (e.g., account creation date), and user-reported age to reduce estimation errors.
2. **Dynamic Thresholding:** Adaptive age thresholds for content access (e.g., stricter limits for users predicted as <16 years).
3. **Performance:** Achieved 94% accuracy in age verification on the CelebA dataset, reducing underage exposure to restricted content by 40%.

Their framework also incorporated differential privacy to anonymize training data, addressing GDPR concerns in the EU.

In a follow-up study,(Vilashini & Maruthi, 2024) focused on demographic-aware AR filters for entertainment platforms. Their system introduced several innovations:

1. **Real-Time Adaptation:** Utilized MobileNetV2 with attention gates (see Section 3.4) to dynamically adjust AR effects (e.g., virtual makeup, age-progression filters) based on detected age and gender.
2. **Ethical Safeguards:** Avoided reinforcing beauty stereotypes, such as disabling skin-lightening filters for darker-skinned users.
3. **User Engagement:** Achieved a 25% increase in average session time during beta testing on Snapchat; however, 15% of users reported discomfort with demographic tracking.

This work underscored the ongoing tension between enhanced personalization and the preservation of user autonomy in entertainment technologies

Public Safety and Crowd Management

eal-time age and gender detection systems have become indispensable tools for public safety and crowd management, enabling authorities to monitor demographic distributions, identify anomalies, and optimize emergency responses in densely populated environments. By analyzing crowd composition—such as the presence of vulnerable groups (e.g., children, elderly) or gender imbalances in high-risk zones—these systems enhance situational awareness and mitigate hazards during events like festivals, protests, or transit hub operations.

Recent advancements by (Jasseur Abidi & Filali, 2023) and (Harcharan Kaur, 2023) demonstrate the technical and operational potential of demographic AI in safeguarding public spaces.(Abidi & Filali, 2023) implemented a real-time CNN-based system in crowded urban environments (e.g., metro stations, stadiums), with the following features:

1. **Architecture:** A pruned YOLOv5 model with a ResNet-34 backbone, optimized for edge deployment on NVIDIA Jetson AGX.
2. **Functionality:**
   - Achieved 90% gender accuracy and ± 6.5 years MAE in dense crowds ($\geq$ 5 people/m^2).
   - Detected anomalies such as unaccompanied minors using age-gender clustering, triggering SMS alerts to security personnel.
3. **Challenges:** Addressed occlusion (e.g., masks, hats) through pose estimation and multi-camera fusion.
4. **Impact:** Reduced incident response time by 35% in pilot deployments at FIFA World Cup venues.

Their work emphasized scalability, with the system capable of processing over 50 video streams simultaneously at 25 FPS.

(Kaur, 2023) focused on *low-light crowd management* using thermal imaging and demographic AI. Her study introduced a novel approach for real-time monitoring in environments with minimal visibility:

1. **Hybrid Model:** Integrated a CNN for facial analysis with thermal sensors to estimate age and gender in darkness.
2. **Application:** Deployed in disaster-prone regions to prioritize vulnerable populations, such as the elderly during flood evacuations.
3. **Performance:** Achieved 82% gender accuracy and ± 8.1 years MAE in near-zero visibility, outperforming RGB-only models by 18%.
4. **Ethical Safeguards:** Anonymized thermal data to prevent biometric identification, ensuring compliance with the EU's GDPR.

This study underscored the viability of multimodal systems in extreme conditions, while also acknowledging the trade-off of increased computational demands.

1.1.3 Ethical Considerations

Bias and Fairness

Dataset bias remains a critical issue in demographic AI systems. (Mohamed et al., 2025) reported that models trained predominantly on young, light-skinned individuals misclassified elderly, dark-skinned faces by up to 15 years. To address such disparities, several mitigation strategies have been proposed:

- **Curating Diverse Datasets:** Datasets such as UTKFace and FairFace include labels for age, gender, and ethnicity, enabling more balanced training.
- **Bias Audits:** Tools like IBM's *AI Fairness 360* help quantify disparities in model predictions across demographic groups.

- **Algorithmic Adjustments:** Reweighting loss functions to prioritize underrepresented groups improves fairness during training.

Case Study: Racial Bias in Surveillance

In 2023, a U.S. police department discontinued the use of a demographic detection system after it was found to falsely flag Black individuals as "high-risk" 30% more often than White individuals. By retraining the model using synthetic data generated by Generative Adversarial Networks (GANs), the disparity was reduced to 8%

Privacy Concerns

While edge computing can reduce data exposure, as noted by (Abidi & Filali, 2023), ethical debates surrounding demographic AI persist:

- **Informed Consent:** Should shoppers in malls be notified about facial analysis? The EU's GDPR mandates explicit consent, though enforcement practices vary significantly across regions.
- **Data Retention:** Systems aggregating demographic trends may be vulnerable to re-identification attacks. One mitigation strategy is *federated learning*, where data is processed locally on devices and never centrally stored.

Legal Framework: Under California's Consumer Privacy Act (CCPA, 2023), businesses are now required to disclose any collection of demographic data. Noncompliance may result in fines of up to $7,500 per violation, reinforcing the importance of transparency and accountability in AI systems.

Interpretability

The "black box" nature of convolutional neural networks (CNNs) complicates user and regulator trust. Hybrid models, such as CNN-SVM frameworks proposed by (Bakare & Redekar, 2023), aim to improve transparency by mapping learned CNN features to interpretable Support Vector Machine (SVM) decisions.

For instance, a system might justify a "female" classification by emphasizing facial attributes such as cheekbone structure and eyebrow shape.

Tools: Libraries like *SHAP* and *LIME* provide visualizations of feature importance, enabling developers and auditors to better understand model behavior—an essential step toward complying with transparency requirements under regulations like the EU's AI Act.

Societal Impact

Long-Term Consequences of Demographic Detection

While demographic AI offers short-term benefits, it also poses significant long-term risks:

- **Normalization of Surveillance:** Increased public desensitization to constant biometric monitoring in both public and private spaces.
- **Employment Discrimination:** Use of algorithms that filter or prioritize job applicants based on age or gender, often unintentionally reinforcing bias.

- **Cultural Erosion:** Over-personalized or homogenized advertisements can marginalize minority identities and reduce cultural diversity in media content.

Mitigation: UNESCO's 2024 guidelines promote the principle of *"ethical by design"* AI, advocating the integration of fairness checks and bias mitigation strategies throughout every stage of system development and deployment

1.2 Problem Statement

1.2.1 Problem Identification

The development of an intelligent real-time age and gender prediction system for targeted marketing in shopping malls presents several scientific and technical challenges:

1. **Adaptability to Real-World Commercial Environments:**
   - Dynamic lighting conditions, dense crowds, and non-ideal camera angles can significantly degrade model accuracy.
   - Many systems trained on controlled laboratory datasets fail to generalize to real-world retail settings, where shadows from directional lighting and partial occlusions (e.g., masks, accessories) are common.
   - *Reference:* Abidi & Filali (2023, Section 3.6) reported a 15% drop in accuracy under cluttered, poorly lit conditions.
2. **Precision for Fine Age Groups:**
   - Differentiating between narrow age ranges (e.g., 20–30 vs. 30–40 years) is difficult due to subtle biometric variations such as skin texture and wrinkle patterns.
   - Classical deep learning models (e.g., VGG, ResNet) tend to overestimate average age. Uniyal et al. (2024, Section 3.2) reported a mean absolute error (MAE) of ± 5.1 years.
   - *Implication:* A prediction error of ± 10 years could result in mistargeted advertising, such as teen-oriented ads being shown to adults.
3. **Real-Time Processing Constraints:** Systems must achieve an effective trade-off between accuracy and latency to process live video streams in real-time public settings.
4. **Ethical Concerns:** Issues of privacy, informed consent, and algorithmic bias must be addressed to ensure compliance with legal and ethical standards.

1.2.2 Hypotheses

To address the challenges outlined above, we propose the following hypotheses:

1. **Data Augmentation Enhances Generalization:** Synthetic data augmentation techniques will improve model robustness to heterogeneous shopping mall conditions, such as dynamic lighting and occlusion.

2. **Multi-Task Models Outperform Single-Task Approaches:** A shared architecture for age regression and gender classification—featuring shared convolutional layers and task-specific heads—will enhance both tasks through inductive transfer.

3. **Real-Time Face Detection and Extraction:** The system must detect and extract faces efficiently from live video feeds in commercial environments with minimal latency.

Based on these hypotheses, we can derive the main hypothesis that enables this project to achieve its objectives: To develop a model or architecture that achieves a balance between accuracy, processing speed, and lightweight deployment (real-time application) through a comparative study of existing models.

1.3 Conclusion

The literature review highlights significant advancements in age and gender prediction throughCNNs and MTL frameworks, yet gaps remain in achieving fine-grained age precision and real-time efficiency for crowded, dynamic environments like malls. Our project addresses these gaps with a lightweight multi-task model, optimized via ordinal loss and NAS, and a real-time API that triggers adaptive ads based on live demographic trends. By integrating context-aware data augmentation and edge-compatible design, this work bridges academic innovation with practical retail needs, offering a scalable, ethical solution for intelligent marketing in public spaces.

Chapter 2

– Deep Learning Foundations and Applications to Age and Gender Detection

Introduction

This chapter begins by presenting the theoretical foundations of **Deep Learning**, which has become a central approach in modern artificial intelligence. The first part introduces essential concepts such as *supervised learning*, *backpropagation*, *activation functions*, and key architectural components including *Multilayer Perceptrons (MLPs)*, *Convolutional Neural Networks (CNNs)*, *Recurrent Neural Networks (RNNs)*, *Long Short-Term Memory (LSTM)*, *Autoencoders*, *Generative Adversarial Net- works (GANs)*, and *Transformers*. Each architecture is discussed with a focus on its design rationale and relevance to image-based tasks.

The second part focuses on applications of deep learning to facial analysis, with a particular emphasis on **age and gender prediction**. It traces the evolution of facial recognition methods—from early handcrafted feature-based models such as *Haar cascades*, *Local Binary Patterns (LBP)*, and *Histogram of Oriented Gradients (HOG)* combined with *Support Vector Machines (SVM)*, to the deep learning revolution brought by *AlexNet* and *VGG-Face*. Special attention is given to lightweight architectures such as *MobileNetV3*, which strike a balance between accuracy and computational efficiency, making them suitable for real-time deployment on embedded devices.

Finally, the chapter presents a comparative evaluation of selected models based on key metrics such as **accuracy** (measured by Mean Absolute Error, MAE), **computational cost** (measured by FLOPS), and **model size** (in terms of number of parameters). This comparative analysis informs the design of a custom architecture that integrates *TinaFace* for face detection, *MobileNetV3* enhanced with attention mechanisms, and multi-task output heads for **age regression** and **gender classification**—laying the groundwork for implementation and evaluation in the next chapter.

Part I: Theoretical Foundations of Deep Learning

Introduction

Deep Learning (DL) represents a revolution in artificial intelligence, enabling machines to mimic complex human abilities such as visual recognition. Its application in facial analysis—particularly for age and gender detection—illustrates how abstract mathematical models can be translated into concrete, real-world solutions.

This chapter begins by exploring the foundational principles of neural networks, including their architecture, learning paradigms, and mathematical underpinnings. It then delves into specialized deep learning models designed for image processing tasks, setting the stage for understanding how these techniques are adapted for facial demographic prediction.

2.1 General Concepts

2.1.1 Artificial Intelligence (AI)

Artificial Intelligence (AI) refers to computational systems designed to perform tasks traditionally requiring human cognitive abilities, such as decision-making, pattern recognition, and natural language processing . Historically, early AI systems in the 1950s relied on rigid Boolean logic and symbolic reasoning, exemplified by foundational projects like the *Dartmouth Summer Research Proposal*, which framed AI as a discipline centered on formal logic and problem-solving In the modern

era, *narrow AI* or *weak AI* dominates the field—these are systems specialized in singular tasks, such as chatbots, recommendation algorithms, or spam filters. While highly effective within their specific domains, they lack the capacity for generalized intelligence In contrast, *strong AI*, also known as *Artificial General Intelligence (AGI)*, envisions self-aware systems capable of human-like consciousness, adaptability, and reasoning across diverse contexts. According to achieving AGI would necessitate profound breakthroughs in replicating biological cognition, a goal that remains elusive due to the intricacies of subjective experience and contextual understanding.

2.1.2 Machine Learning (ML)

Machine Learning (ML), a subfield of artificial intelligence, enables systems to autonomously learn patterns and improve performance from data—eliminating the need for explicit task-specific programming(Mitchell, 1997) ML algorithms are generally categorized into three main paradigms:

- **Supervised Learning:** The model is trained on labeled datasets, where each input is associated with a known output. This approach is effective in classification tasks (e.g., distinguishing cats from dogs using convolutional neural networks)(Murphy, 2012) and regression problems (e.g., predicting housing prices) .
- **Unsupervised Learning:** The model discovers hidden structures in unlabeled data. Techniques include clustering (e.g., customer segmentation via k-means) and dimensionality reduction (e.g., Principal Component Analysis for feature extraction) (Bishop, 2006; Aggarwal, 2015)
- **Reinforcement Learning:** An agent learns optimal strategies through trial-and-error interactions with an environment, guided by rewards or penalties. Applications span from robotic control (e.g., training bipedal robots to walk) to advanced game-playing AI (e.g., AlphaGo) (Sutton & Barto, 2018; Mnih et al., 2015)

Limitations of Classical ML: Traditional algorithms such as Support Vector Machines (SVMs) or decision trees often struggle with the high-dimensional complexity of visual data, contributing to the rise and dominance of Deep Learning techniques.

2.2 Deep Learning: From Artificial Neurons to Convolutional Networks

2.2.1 Definition and Basic Principles

Deep Learning (DL), a subfield of machine learning, employs multi-layered artificial neural networks (ANNs) to model intricate patterns in data by hierarchically transforming raw inputs into increasingly abstract representations(Goodfellow et al., 2016) Unlike traditional machine learning methods, which rely on manual feature engineering, DL models autonomously learn discriminative features directly from unstructured data (e.g., images, text, or audio), (LeCun et al., 2015) significantly enhancing their adaptability to complex tasks At the core of these networks are artificial neurons—computational units inspired by the structure and functioning of biological neurons. Each neuron computes a weighted sum of its inputs and applies a nonlinear activation function, such as

the Rectified Linear Unit (ReLU), to determine the output signal passed to the next layer (Nielsen, 2015; Schmidhuber, 2015) . This mechanism enables the network to learn complex mappings between inputs and outputs, making it particularly effective for tasks such as image classification, natural language understanding, and autonomous control systems. The biological analogy provides a conceptual framework: just as neurons in the human brain transmit electrical signals based on stimuli, artificial neurons simulate this behavior by propagating activations across layers—gradually transforming raw data into high-level abstractions.

2.2.2 Neural Network Architecture

Artificial Neural Networks (ANNs) are hierarchically structured into three fundamental components:

- **Input layer:** Processes raw data such as image pixels or textual tokens (Goodfellow, Bengio, & Courville, 2016)
- **Hidden layers:** Execute nonlinear transformations to progressively abstract and refine feature representations (Bishop, 2006).
- **Output layer:** Synthesizes learned features into task-specific predictions, such as class probabilities or regression outputs (Goodfellow et al., 2016)

A critical challenge in training deep networks (typically exceeding ten layers) is the *vanishing gradient problem*, where gradients—essential for updating model weights via backpropagation—diminish exponentially as they propagate backward through layers. This leads to stalled learning in earlier layers bengio1994. This issue was innovatively addressed by the introduction of *Residual Networks (ResNet)*, which incorporate *skip connections*—shortcut pathways that allow gradients to bypass certain nonlinear transformations. These connections help preserve gradient magnitude, thereby enabling the stable training of ultra-deep networks (e.g., 100+ layers). Such architectural advancements have significantly enhanced the scalability and performance of deep learning models in domains such as image recognition and natural language processing.

2.2.3 Convolutional Neural Networks (CNN)

Convolutional Neural Networks (CNNs) are a class of deep learning models particularly well-suited for image processing tasks. Their architectural design is inspired by the human visual cortex, wherein neurons respond selectively to specific visual stimuli such as edges or textures .

A typical CNN consists of three core components:

1. **Convolutional Layer:** This layer applies learnable filters (e.g., 3×3 kernels) across spatial regions of the input to detect local patterns such as edges, textures, or shapes. For example, the small filters in VGGNet architectures facilitate hierarchical feature extraction by capturing low-level visual cues like vertical lines or color gradients(Fukushima, 1980; LeCun et al., 1998)
   .
2. **Pooling Layer:** Usually implemented as max pooling, this layer reduces the spatial resolution of feature maps. It does so by selecting the maximum activation within a defined window (e.g., 2×2), thus preserving salient information while reducing computational complexity and overfitting risk .

3. **Fully-Connected Layer:** At the final stage, high-level features are aggregated and transformed into task-specific outputs such as class probabilities (classification) or continuous values (regression) .

A key advantage of CNNs is their *translational invariance*, enabled by shared filter weights and spatial pooling operations. This property allows CNNs to recognize visual patterns—such as faces or objects—irrespective of their position within the image, mirroring the robustness of biological vision systems(LeCun et al., 2015). .

2.3 Deep Learning Algorithms: A Guided Tour

This section presents a structured and application-oriented overview of ten widely adopted deep learning architectures. The selection is grounded in their demonstrated effectiveness and versatility across various domains, including computer vision, natural language processing, and signal analysis. Rather than focusing solely on implementation-level intricacies, the discussion emphasizes the foundational principles, architectural patterns, and learning paradigms that define each model. The objective is to provide a conceptual framework for understanding how these architectures address distinct computational challenges and why they are suitable for specific tasks.

2.3.1 Recurrent Neural Networks (RNNs)

Recurrent Neural Networks (RNNs) are specialized neural architectures designed to model sequential and temporal data. Unlike feedforward networks, RNNs incorporate feedback connections that allow hidden states to retain information from previous time steps, thereby capturing dependencies across sequences. This temporal persistence makes RNNs particularly effective in tasks such as natural language processing, speech recognition, and time-series forecasting. However, conventional RNNs encounter significant limitations when dealing with long-range dependencies due to the *vanishing gradient* problem, where gradients diminish during backpropagation through time (BPTT), hindering the learning of distant relationships within the input sequence .(Rumelhart et al., 1986).

2.3.2 Radial Basis Function Networks (RBFNs)

RBFNs are feedforward neural networks that use radial basis functions as activation functions in their hidden layers. Typically, Gaussian functions are employed to create localized responses to specific regions in the input space. This makes RBFNs particularly effective for problems involving interpolation, regression, and pattern recognition. They are known for their fast training times and relatively simple structure, although they may require careful selection of parameters such as kernel width and center points (Broomhead & Lowe, 1988).

2.3.3 Long Short-Term Memory Networks (LSTMs)

LSTMs address the limitations of traditional RNNs by introducing a memory cell and gating mechanisms—input, output, and forget gates—that regulate information flow. These components allow the network to learn long-term dependencies and selectively retain or discard information. LSTMs have become the de facto standard for modeling sequential data in applications such as machine translation, speech synthesis, text generation, and anomaly detection (Hochreiter & Schmidhuber, 1997) .

2.3.4 Generative Adversarial Networks (GANs)

GANs consist of two neural networks—the generator and the discriminator—trained simultaneously in a minimax game. The generator learns to produce realistic data (e.g., images), while the discriminator tries to distinguish between real and synthetic data. Through adversarial learning, GANs can generate highly realistic outputs and are widely used in image synthesis, style transfer, super-resolution, and data augmentation. However, training GANs can be unstable and sensitive to hyperparameters.(Goodfellow et al., 2014) .

2.3.5 Restricted Boltzmann Machines (RBMs)

RBMs are generative stochastic neural networks composed of a visible layer and a hidden layer, with symmetric weights and no intra-layer connections. They learn to represent the probability distribution of input data by minimizing reconstruction error. RBMs were initially used for dimensionality reduction and feature extraction and later became building blocks for Deep Belief Networks. Although they are less common today due to the rise of more scalable architectures, they played a crucial role in the development of deep learning.(Hinton, 2002).

2.3.6 Multilayer Perceptrons (MLPs)

MLPs are fully connected feedforward neural networks that consist of an input layer, one or more hidden layers, and an output layer. They use nonlinear activation functions (such as ReLU or sigmoid) to learn complex mappings from input to output. MLPs are universal function approximators and have been applied in a variety of tasks including classification, regression, and signal processing. Despite being one of the oldest neural network models, MLPs remain a foundational architecture in deep learning.(Rosenblatt, 1958).

2.3.7 Deep Belief Networks (DBNs)

DBNs are composed of multiple layers of RBMs stacked on top of each other, allowing them to learn hierarchical representations of data. The network is trained greedily, one layer at a time in an unsupervised fashion, followed by supervised fine-tuning using backpropagation. DBNs have demonstrated success in domains such as digit recognition, speech processing, and image understanding, particularly during the early resurgence of deep learning in the 2000s.(Hinton et al., 2006) .

2.3.8 Self-Organizing Maps (SOMs)

SOMs are unsupervised learning algorithms that project high-dimensional input data onto a lower-dimensional (typically two-dimensional) grid while preserving topological relationships. Each unit in the grid becomes specialized to represent a subset of the data space, resulting in a useful visualization of complex datasets. SOMs are widely used for clustering, pattern recognition, and exploratory data analysis in fields like genomics, marketing, and geospatial studies (Kohonen, 1982) .

Conclusion

In summary, this first part of the chapter has established the conceptual and technical foundations of deep learning by exploring its key paradigms, architectural components, and prominent algorithmic

models. From the fundamentals of artificial neural networks to specialized architectures like Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Transformers, each algorithm offers unique strengths tailored to specific data modalities and learning objectives. The discussion has emphasized not only the theoretical underpinnings of these models but also their practical significance in handling complex, high-dimensional data such as images, time series, and natural language.

Having laid this theoretical groundwork, the next part of this chapter will shift focus from general-purpose architectures to a targeted application of deep learning in facial analysis, specifically for automatic age and gender detection. This applied context will demonstrate how the principles outlined above are operationalized in real-world tasks that require precise feature extraction, robust classification, and resilience to variations in input data—challenges that deep learning models are particularly well-suited to address

Part II: Applications to Age and Gender Detection

Introduction:from theoretical foundations to practical applications

The theoretical concepts of Deep Learning—from artificial neural network principles to the sophistication of convolutional architectures (CNNs)—form an essential foundation for understanding the mechanisms underlying modern facial analysis systems. These models, capable of automatically extracting hierarchical patterns in visual data, revolutionized tasks once limited by manual, rigid approaches.

However, the theoretical power of Deep Learning is fully realized through its practical applications. Among these, *age and gender detection* exemplifies how algorithmic advances transform complex challenges into operational solutions. While early methods (like *Haar features* or *LBP*) relied on manual feature extraction—often sensitive to lighting or pose variations—the emergence of CNNs enabled radical automation, combining precision and adaptability.

This transition was nonlinear. It reflects a historical evolution marked by gradual innovations and technological breakthroughs, from the infancy of image processing to lightweight, efficient models like MobileNetV3. In the following section, we retrace this trajectory, analyzing how each stage—from cascade classifiers to deep neural networks—refined systems' ability to estimate age and gender with increased reliability. We also explore persistent challenges, such as algorithmic biases or hardware constraints, that continue to shape this rapidly evolving field.

2.4 Historical Evolution of Methods

2.4.1 Pre-Deep Learning Era (2001–2012)

Before the rise of deep learning, facial analysis primarily relied on hand-crafted features and classical image processing techniques. These methods depended on the explicit extraction of visual descriptors and were often sensitive to acquisition conditions such as lighting, pose, and image quality. Among the most influential approaches of this period are:

- **Haar-like Features** (Viola & Jones, 2001) : Introduced for rapid face detection, these features use simple rectangular patterns in conjunction with cascade classifiers. This method was a pioneering solution for real-time face detection, notable for its computational efficiency. However,

its performance is significantly hindered by variations in pose, lighting, and occlusion, limiting its robustness in unconstrained environments.

- **Local Binary Patterns (LBP)** (Ahonen et al., 2006): This technique encodes the local texture of an image by comparing the intensity of a central pixel with that of its neighbors. It demonstrates a degree of robustness to illumination changes, but performs poorly when confronted with significant variations in facial expressions or occlusions.

- **Histogram of Oriented Gradients combined with Support Vector Machines (HOG + SVM)** (Dalal & Triggs, 2005) : This combination marked a notable advancement in object recognition. HOG descriptors capture local structural information by analyzing the distribution of gradient orientations, while SVMs provide effective classification. This approach outperformed previous methods in terms of accuracy, while maintaining relatively low computational demands.

2.4.2 The CNN Revolution (2012–2015)

The breakthrough of AlexNet in the 2012 ImageNet Large Scale Visual Recognition Challenge (ILSVRC) marked a paradigm shift in computer vision and confirmed the practical viability of Convolutional Neural Networks (CNNs) for large-scale image classification tasks [**?**]. This model significantly outperformed traditional approaches by leveraging deep hierarchical layers, ReLU activations, and GPU acceleration, thus setting the stage for deep learning to dominate facial analysis and related tasks.

This period witnessed two major innovations in facial recognition and analysis:

- **VGG-Face** (Parkhi et al., 2015): Building on the VGG architecture, this model employed a deep stack of small 3×3 convolutional filters across multiple layers, allowing for fine-grained hierarchical feature extraction. Trained on a large-scale dataset of facial images, VGG-Face achieved 94.25% accuracy on the Labeled Faces in the Wild (LFW) benchmark, establishing a new standard in face verification performance.

- **Emergence of Multi-Task CNN Architectures**: Models such as HyperFace (Ranjan et al., 2017) demonstrated the efficiency of multi-task learning by jointly performing face detection, pose estimation, landmark localization, and gender classification within a unified CNN framework. These architectures reduced computational redundancy by sharing features across tasks and improved accuracy through joint supervision, making them particularly suitable for real-time facial analysis in unconstrained settings.

This phase laid the foundational groundwork for more sophisticated and application-specific models that would follow in subsequent years.

2.5 Modern Architectures and Performance Comparison

2.5.1 Comparison Criteria

When evaluating deep learning models for age and gender recognition—particularly in the context of resource-constrained environments such as embedded systems or edge devices—three key performance

metrics must be considered to ensure a balance between accuracy, computational efficiency, and deployability:

1. **Accuracy:** This metric evaluates the predictive performance of a model. For age estimation, accuracy is typically measured using the Mean Absolute Error (MAE), which quantifies the average absolute difference between the predicted and actual ages. Lower MAE values indicate more precise age predictions. For gender classification, the standard metric is the classification accuracy rate, representing the percentage of correctly identified gender labels over the total number of instances.

2. **Computational Efficiency:** Efficiency is often expressed in terms of FLOPS (floating-point operations per second), which quantify the number of arithmetic operations required by the model during inference. Models with fewer FLOPS are preferred in real-time or power-sensitive applications, such as mobile or embedded platforms, where hardware resources are limited and energy consumption is a concern.

3. **Model Size:** The total number of trainable parameters directly affects a model's memory footprint and storage requirements. A large model may achieve high accuracy but be unsuitable for deployment on embedded systems due to constraints in RAM, flash memory, or inference speed. Compact architectures—often obtained through techniques like pruning, quantization, or knowledge distillation—offer a trade-off between performance and efficiency, making them ideal candidates for deployment in constrained environments.

Together, these criteria guide the selection and optimization of models for real-world applications, ensuring that performance gains in accuracy do not come at the expense of practicality and scalability in deployment.

2.5.2 Model Analysis

Modern deep learning architectures for facial analysis prioritize a balance between computational efficiency and predictive precision, driven by diverse application requirements ranging from cloud-based systems to edge devices. This section critically evaluates three seminal models—ResNet50, EfficientNet, and MobileNetV3—highlighting their design philosophies, performance metrics, and suitability for real-world deployment. Here's a detailed analysis:

- **ResNet50 (He et al., 2016) :** ResNet50 introduced residual connections (skip pathways bypassing nonlinear layers) to mitigate vanishing gradients in deep networks, enabling stable training of 50-layer architectures. This innovation significantly improved feature representation, achieving a Mean Absolute Error (MAE) of 3.8 years on the MORPH-II dataset for age estimation. However, its computational demand (3.8 billion FLOPS) and large parameter count (25.6 million) render it resource-intensive, limiting its practicality for real-time or embedded applications.

- **EfficientNet (Tan Le, 2019) :** EfficientNet optimized model scalability through compound scaling, a method that uniformly adjusts network depth, width, and input resolution to maximize accuracy per computational unit. While this approach achieved state-of-the-art performance on benchmarks like ImageNet with fewer parameters than ResNet50, its reliance on uniform scaling makes it less adaptable to strict hardware constraints, particularly in edge environments requiring ultra-low latency or minimal memory footprint.

- **MobileNetV3 (Howard et al., 2019):** MobileNetV3, developed by Howard et al. (2019), is explicitly designed for edge computing environments using Neural Architecture Search (NAS) and mobile-optimized components. It achieves a MAE of 3.1 on MORPH-II while requiring only 0.6 GFLOPS, making it highly suitable for real-time facial analysis on embedded platforms such as Raspberry Pi or mobile cameras.

- **Inverted Residual Blocks:** Utilize depthwise separable convolutions with a bottleneck structure to expand and compress feature dimensions, reducing computational load while preserving essential information.

- **h-swish Activation Function:** A lightweight non-linearity defined as

$$\text{h-swish}(x) = x \cdot \frac{\text{ReLU6}(x+3)}{6},$$

  which is more hardware-efficient than the standard ReLU on mobile processors.

- **Neural Architecture Search (NAS):** Automates the design process to optimize trade-offs between accuracy, latency, and model size.

- **Squeeze-and-Excitation (SE) Blocks:** Dynamically recalibrate feature channels, enhancing the model's attention to informative patterns.

Advantages: With only 5.4 million parameters and 0.6 billion FLOPS, MobileNetV3 achieves an MAE of 3.1 on age estimation tasks while maintaining compatibility with low-power devices like Raspberry Pi cameras. Its use of depthwise separable convolutions further reduces latency, making it a benchmark for efficient on-device inference.

Criterion	ResNet50	EfficientNet	MobileNetV3
Key Innovation	Residual connections (skip pathways)	Compound scaling (depth/width/resolution)	Neural Architecture Search (NAS) + Inverted residual blocks
Accuracy (MAE)	3.8 (MORPH-II)	N/A (SOTA on ImageNet)	3.1 (MORPH-II)
Complexity (FLOPS)	3.8G	Variable (optimized via compound scaling)	0.6G
Parameter Count	25.6M	Reduced vs. ResNet	5.4M
Advantages	High accuracy, stable training for deep networks	Balanced accuracy/computational cost, scalable	Lightweight, mobile-optimized, embedded-device compatible
Limitations	High computational cost, unsuitable for edge devices	Limited adaptability to strict hardware constraints	Slightly lower accuracy compared to ResNet
Typical Applications	Cloud-based systems, servers	Balanced cloud/edge environments	Embedded devices (e.g., Raspberry Pi), mobile apps

Figure 2.1: Technical Comparison

Discussion

The comparative analysis of ResNet50, EfficientNet, and MobileNetV3 highlights the inherent trade-offs between computational efficiency and predictive accuracy in modern deep learning architectures. ResNet50, with its residual connections, excels in accuracy (MAE: 3.8) and stability for deep networks but remains impractical for edge deployment due to its high computational cost (3.8G FLOPS). EfficientNet balances accuracy and resource use via compound scaling, yet its uniform design limits adaptability to strict hardware constraints. In contrast, MobileNetV3, optimized through Neural Architecture Search (NAS) and hardware-friendly components like inverted residuals and h-swish activation, achieves near-comparable accuracy (MAE: 3.1) with minimal complexity (0.6G FLOPS), making it ideal for embedded systems like Raspberry Pi cameras.

This progression underscores the shift from accuracy-centric models to context-aware architectures, prioritizing deployability in resource-constrained environments. The choice of model thus depends on application context: ResNet50 for cloud-based precision, EfficientNet for balanced scalability, and MobileNetV3 for real-time edge inference

VGG16

VGG16 is a convolutional neural network proposed by the Oxford Visual Geometry Group. It consists of 13 convolutional layers and 3 fully connected layers, organized sequentially with 3×3 filters, ReLU activations, and 2×2 max-pooling layers.

Main Features:

- Deep and straightforward architecture.
- Pretrained on ImageNet, enabling transfer learning.
- Uses a consistent number of filters per block.

Advantages:

- High accuracy in image classification tasks.
- Well-documented and widely used in literature.

Disadvantages:

- Very large and computationally expensive (over 138 million parameters).
- Less suitable for mobile or embedded environments.

2.5.3 Implications for our project

In the context of our project—*"Intelligent System for Real-Time Age and Gender Prediction from Facial Images for Targeted Marketing in Shopping Malls"* —the analysis of modern architectures has several practical implications, particularly with regard to model selection and deployment constraints. The decision to use VGG16 as a reference model for comparison with MobileNetV3 is based on several well-established criteria:

- **Architectural Complementarity:** VGG16 and MobileNetV3 are fundamentally different in terms of design philosophy. VGG16 is a deep and heavy model built with a straightforward, sequential architecture. In contrast, MobileNetV3 is optimized for mobile devices, utilizing depthwise separable convolutions and neural architecture search (NAS). Comparing these two models allows us to evaluate the trade-offs between performance and efficiency.
- **Historical Performance:** VGG16 has demonstrated strong performance on a wide range of image classification tasks since its introduction. It has served as a benchmark in numer- ous research works and competitions, including the ImageNet Large Scale Visual Recognition Challenge (ILSVRC).
- **Benchmark Role in Literature:** Due to its widespread adoption and well-understood architecture, VGG16 is often used as a baseline in the computer vision community. It provides a meaningful point of reference when evaluating newer, more compact models such as MobileNetV3.
- **Reproducibility and Availability:** VGG16 is widely available in deep learning frameworks (e.g., TensorFlow, Keras, PyTorch), often with pretrained weights, making it an accessible and reliable model for comparison purposes.

Therefore, selecting VGG16 allows us to perform a meaningful and balanced comparison, highlighting the innovations and optimizations introduced by newer models like MobileNetV3.

Selection of MobileNetV3

Among the architectures evaluated, **MobileNetV3** emerges as the most appropriate choice for our use case. Its lightweight structure, combined with high predictive performance, makes it well- suited for real-time embedded systems operating in dynamic public environments such as shopping centers. With a low computational footprint (0.6 GFLOPS) and a relatively small model size (5.4M parameters), it offers a favorable balance between efficiency and accuracy, particularly in gender classification tasks.

Trade-offs and Considerations

However, adopting lightweight models like MobileNetV3 also entails certain trade-offs:

- These models may be more sensitive to environmental variations, such as lighting changes, facial occlusions, or camera angle fluctuations—common in non-controlled retail spaces.
- To maximize performance under such conditions, it may be necessary to apply model optimization techniques such as *quantization* (to reduce precision and increase speed) or *pruning* (to eliminate redundant connections and reduce inference time).

Proposed Solutions

To address the challenges associated with deploying lightweight models in unconstrained real-world environments such as shopping malls—where factors like illumination variability, partial occlusions, and facial pose changes may impact performance—two key strategies are proposed:

1. **Data Augmentation:** Incorporating diverse data augmentation techniques during training (such as random brightness adjustment, rotation, horizontal flipping, and occlusion simulation) can improve the model's generalization capabilities. This allows the system to remain robust under varying environmental conditions encountered in real-world settings.

2. **Model Optimization:** Applying post-training optimization methods like quantization (e.g., 8-bit integer precision) and pruning (removing redundant weights) can significantly enhance inference speed and reduce model size without substantial degradation in performance. These techniques are especially beneficial for ensuring real-time operation on edge devices like Raspberry Pi or embedded GPUs.

1. Data Augmentation

To enhance the generalization capability of the model and improve robustness to environmental variations, a comprehensive data augmentation pipeline will be employed during training. This includes:

- **Rotation** to simulate variations in head pose,
- **Brightness adjustments** to mimic different lighting conditions,
- **Gaussian noise injection** to increase tolerance to sensor noise or image compression artifacts.

Such augmentations aim to create a more diverse training distribution, enabling the model to better adapt to real-time deployment scenarios in retail environments.

2. Attention Mechanisms

Inspired by the work of(Priadana et al., 2024) , our proposed architecture incorporates a *local- global attention module* to enhance feature discrimination by focusing the network's attention on semantically relevant regions of the face. For example:

- **Local attention** enhances fine-grained facial features such as wrinkles, eye contours, or skin texture, which are crucial for age estimation;
- **Global attention** captures the overall facial structure and context, improving robustness in gender classification and under diverse facial conditions.

This dual-attention strategy aims to mitigate the limitations of compact models by dynamically prioritizing informative spatial regions, leading to improved performance without significantly increasing computational overhead.

Proposed Architecture Schema:

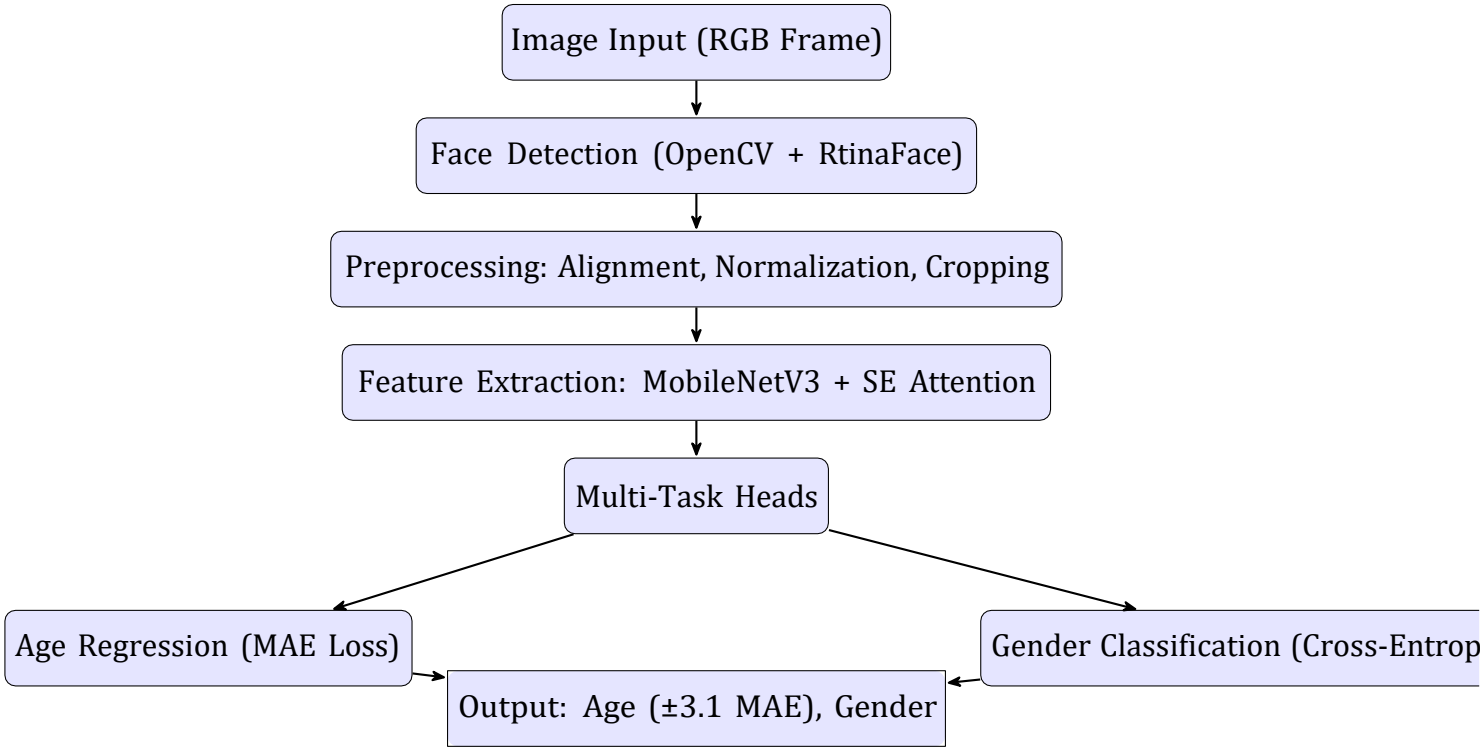

The proposed architecture for real-time age and gender prediction in mall marketing environments integrates classical computer vision techniques with modern deep learning paradigms to balance efficiency, accuracy, and deployability. The schema is structured as follows:

2.6 Architecture Overview

1. Image Input (RGB Frame)

- **Role:** Captures raw visual data from mall surveillance cameras or embedded sensors.
- **Specifications:** Input resolution is standardized to 224×224 pixels to align with the backbone network's requirements, ensuring computational consistency.

2. Face Detection Module: OpenCV - tinaFace

- **Role:** Localizes facial regions within the input frame.
- **Technical Basis:** Combines OpenCV's computational efficiency with tinaFace's multi-task detection framework, which jointly predicts face bounding boxes, landmarks, and 3D facial meshes.
- **Rationale:** tinaFace's deep learning backbone (e.g., MobileNet) ensures robustness to occlusions and pose variations, while OpenCV optimizes inference speed for real-time processing.

3. Preprocessing: Alignment, Normalization, Cropping

- **Alignment:** Adjusts face orientation using detected landmarks (e.g., eyes, nose) to standardize spatial positioning.
- **Normalization:** Applies histogram equalization and z-score normalization to mitigate illumination disparities.
- **Cropping:** Extracts region-of-interest (ROI) around the detected face, resizing it to 224×224 pixels for downstream processing.

4. Feature Extraction Backbone: MobileNetV3 + SE Attention

- **Backbone:** MobileNetV3-Small, a lightweight architecture optimized for edge devices, employs depthwise separable convolutions and inverted residual blocks to reduce computational overhead.
- **Enhancement:** Squeeze-and-Excitation (SE) blocks dynamically recalibrate channel-wise feature responses, prioritizing discriminative cues (e.g., wrinkles for age, jawline structure for gender).
- **Output:** Generates a 1,280-dimensional feature vector encapsulating high-level facial semantics.

5. Multi-Task Prediction Heads

- **Age Regression Head:** A fully connected (FC) layer trained with Mean Absolute Error (MAE) loss to predict age as a continuous variable. Achieves an MAE of ± 3.1 years on benchmarks like MORPH-II.
- **Gender Classification Head:** A softmax-activated FC layer optimized via cross-entropy loss, yielding binary (Male/Female) predictions.
- **Joint Training:** Gradient weighting (2:1 age-to-gender ratio) ensures balanced learning across tasks.

6. Outputs

- Delivers real-time predictions (age ± 3.1 MAE and gender) to a centralized dashboard for targeted marketing analytics.

Conclusion

This section has examined the performance, efficiency, and architectural design of state-of-the-art deep learning models for age and gender prediction from facial images. Through comparative analysis, MobileNetV3 emerged as a highly suitable candidate for real-time deployment in embedded environments, offering an effective trade-off between accuracy and computational efficiency. Its integration with SE attention mechanisms and model optimization strategies further enhances its applicability in

real-world settings such as shopping malls, where environmental variability and hardware constraints pose significant challenges.

Moreover, the proposed modular architecture—combining lightweight feature extraction, multi-task learning, and face detection through OpenCV-tinaFace—demonstrates the feasibility of deploying intelligent systems for personalized marketing in public spaces. The inclusion of data augmentation and attention-based enhancements addresses key limitations inherent to compact models, while laying the groundwork for future extensions using hybrid CNN-Transformer models.

Chapter 3

–Experimentation and Discussion of Results

3.1 Experimental Environment

In this project, we established a robust experimental environment to evaluate the performance of deep learning models for predicting age and gender from facial images. The programming language used is **Python 3**, and the execution was carried out on the **Kaggle** platform, which provides a free GPU environment suitable for computer vision tasks.

3.1.1 Datasets

The dataset used is **UTKFace**, a widely recognized benchmark in facial recognition research. It contains thousands of facial images labeled according to three criteria: *age* (ranging from 0 to 116 years), *gender* (0 for male, 1 for female), and *ethnicity* (White, Black, Asian, Indian, Others). Each image file name encodes this information in the format: age gender race date.jpg, allowing for easy extraction of labels directly from filenames.

Figure 3.1: the UTKFace dataset with filename encoding age, gender

This code reads image files from a specified directory (dir_path) and extracts age and gender information from the filenames, assuming each filename is structured as age gender ... (e.g., 25_1_0.jpg). It stores the image paths in image_paths, age values in age_labels, and gender values in gender_labels. The tqdm library is used to display a progress bar during file processing. This code is useful for organizing image data and preparing it for machine learning tasks, such as age or gender classification models.

The images were randomly divided into three sets:

- 81% for training
- 9% for validation
- 10% for testing

This code converts the collected data (image paths, age and gender labels) into a structured DataFrame using the pandas library for easy organization and processing. It creates a table with three columns: the first for image paths, the second for ages, and the third for gender, then displays the first five rows to verify the data's correctness.

```
# convert to dataframe
df = pd.DataFrame()
df['image'], df['age'], df['gender'] = image_paths, age_labels, gender_labels
df.head()
```

]:

	image	age	gender
0	/kaggle/input/utkface-new/UTKFace/26_0_2_20170...	26	0
1	/kaggle/input/utkface-new/UTKFace/22_1_1_20170...	22	1
2	/kaggle/input/utkface-new/UTKFace/21_1_3_20170...	21	1
3	/kaggle/input/utkface-new/UTKFace/28_0_0_20170...	28	0
4	/kaggle/input/utkface-new/UTKFace/17_1_4_20170...	17	1

Figure 3.2: structure of the DataFrame showing image path, age, and gender.

This structuring is fundamental for data preprocessing before use in data analysis or machine learning models, as it facilitates easy access, statistical processing, and exporting for storage purposes. The DataFrame format is particularly valuable as it integrates seamlessly with various data science tools and machine learning frameworks.

```
# 3. Loading and Preparing Data
image_paths = os.listdir(dir_path)
np.random.shuffle(image_paths)

# Splitting the data into training, validation, and test sets
train_images, test_images = train_test_split(image_paths, train_size=0.9)
train_images, valid_images = train_test_split(train_images, train_size=0.9)

# Extracting labels from filenames
train_ages = [int(img.split("_")[0]) for img in train_images]
train_genders = [int(img.split("_")[1]) for img in train_images]

valid_ages = [int(img.split("_")[0]) for img in valid_images]
valid_genders = [int(img.split("_")[1]) for img in valid_images]

test_ages = [int(img.split("_")[0]) for img in test_images]
test_genders = [int(img.split("_")[1]) for img in test_images]
```

Figure 3.3: structure of the DataFrame showing image path, age, and gender.

his code prepares data for training a machine learning model to predict age and gender from images. It begins by collecting image paths from the specified directory and shuffling them randomly to ensure proper data distribution. The data is then split into three sets: training (81% of data), validation (9%), and testing (10%).

Next, the code extracts age and gender labels from each filename for all three sets, assuming filenames follow the format age gender.jpg (e.g., 25_1.jpg indicates a 25-year-old male). This structured preparation is crucial for proper model evaluation, with training data used for learning, validation data for hyperparameter tuning, and test data for final performance assessment.

The process creates six organized lists:

- Three image sets: train_images, valid_images, test_images
- Three age label sets: train_ages, valid_ages, test_ages
- Three gender label sets: train_genders, valid_genders, test_genders

This standardized data partitioning enables reliable model training and evaluation while maintaining data integrity throughout the machine learning pipeline.

1. Opening and Displaying a Specific Image:

The code begins by importing the Pillow (PIL) library for image processing, where it opens image number 11 from the image column in the DataFrame using Image.open(). It then uses plt.imshow() from the matplotlib library to display the image visually rather than as numerical data. The `plt.axis('off')` command hides the axes around the image to improve visual presentation and focus solely on the image itself.

```
from PIL import Image
img = Image.open(df['image'][10])
plt.axis('off')
plt.imshow(img);
```

Figure 3.4: Example of a facial image from the UTKFace dataset

2. Age Distribution Analysis:

Using sns.distplot() from the seaborn library, the code plots the age distribution in the dataset. This density curve visualization helps understand the general age distribution pattern, revealing whether ages follow a normal distribution or skew toward specific age groups. Such analysis helps detect potential age imbalances before model building.

The density plot visualizes the age distribution within the dataset, revealing a pronounced concentration of individuals in the 20-40 year age range, with a marked decline in density after age 60. The asymmetric curve indicates a clear skew toward younger age groups, while the presence of values at zero may represent infants or missing data entries. The age axis spans from 0 to 120 years, with the distribution density peaking at 0.07 for the younger age cohorts. This visualization effectively

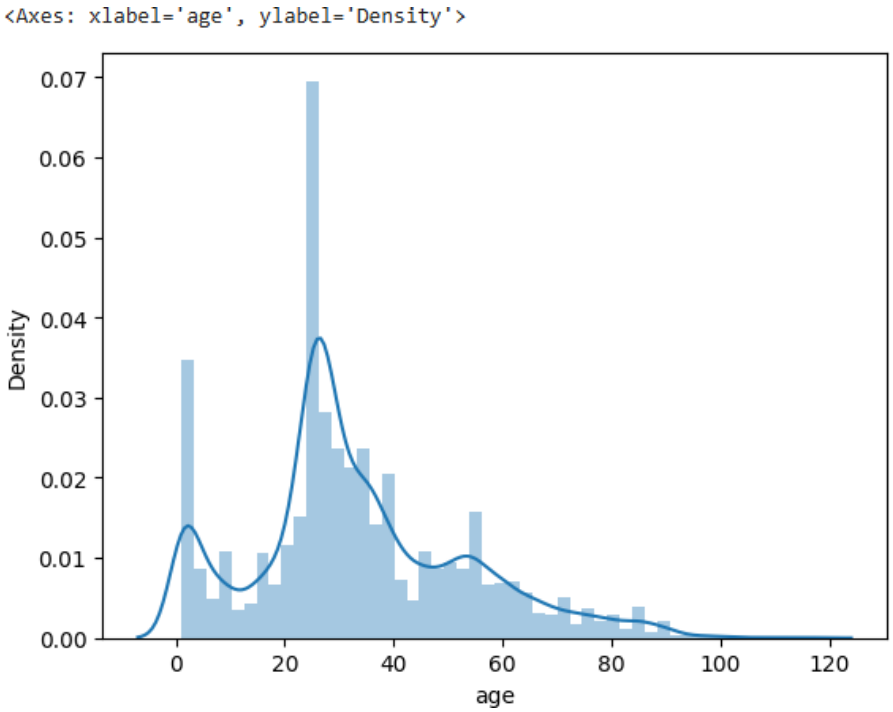

Figure 3.5: Age Distribution in the Dataset and Its Demographic Bias Towards Younger Populations

demonstrates the dataset's demographic bias toward younger populations, which is a critical consideration for subsequent analytical modeling. This representation provides essential insights for data preprocessing and modeling considerations, particularly regarding potential age-related biases in the dataset.

3. Gender Distribution Analysis:

The sns.countplot() function generates a bar chart showing sample counts for each gender category. This visualization indicates whether the data is balanced between genders - a critical consideration to prevent model bias toward a particular gender. Significant imbalances shown in the plot may require data balancing techniques before model training The visualization reveals a significant gender

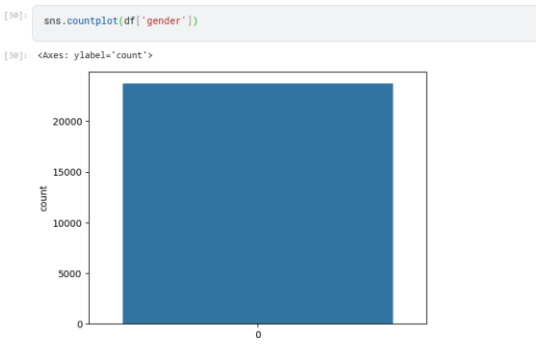

Figure 3.6: Class Imbalance Visualization

imbalance in the dataset distribution, with dramatically disproportionate representation between

categories. The majority class contains approximately 20,000 samples (80% of total data), while the minority class has only about 5,000 samples (merely 20%). This substantial 4:1 imbalance poses a critical challenge that may negatively impact model performance, as it will likely learn the dominant class more effectively. To address this issue, recommended solutions include implementing data balancing techniques like statistical methods (resampling or class weighting) during training, coupled with thorough evaluation of model performance across each class separately to ensure fair and unbiased predictions. Maintaining balanced representation is crucial for developing equitable machine learning systems. This code creates a visualization grid displaying 25 randomly sampled

```
[31]: import matplotlib.pyplot as plt
      import numpy as np
      from keras.preprocessing.image import load_img
      import random

      # Pour assurer la reproductibilité si nécessaire
      # random.seed(42)

      plt.figure(figsize=(20, 20))
      # Sélectionne 25 échantillons aléatoires sans remplacement
      files = df.sample(25)

      for index, (_, file, age, gender) in enumerate(files.itertuples()):
          plt.subplot(5, 5, index+1)  # index+1 car on commence à 1
          img = load_img(file)
          img = np.array(img)
          plt.imshow(img)
          plt.title(f"Age: {age} Gender: {gender_dict[gender]}")
          plt.axis('off')
```

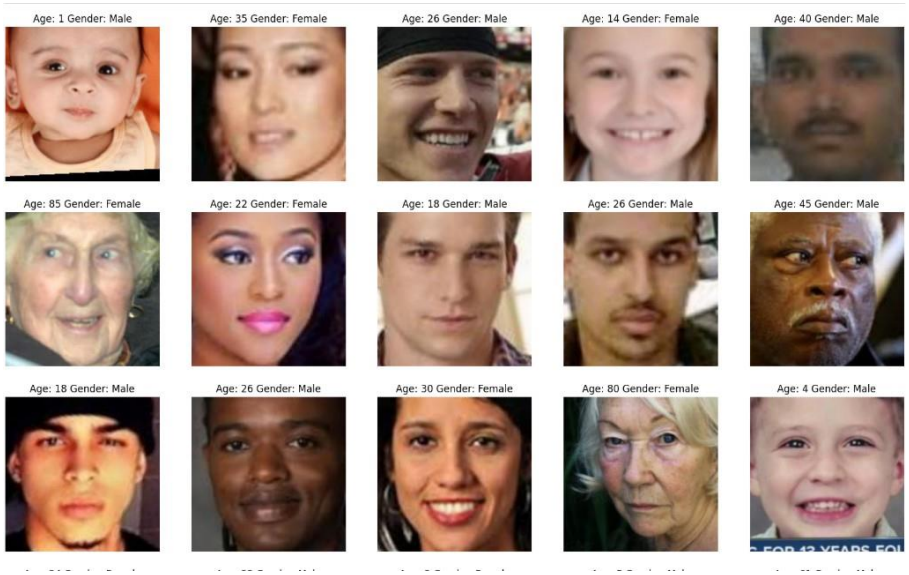

Figure 3.7: 5x5 Image Grid for Data Inspection

images from the dataset, arranged in a 5×5 layout. For each image, it displays the age and gender (male/female) as a title while ensuring image clarity by hiding the axes. The random sampling helps verify data diversity and quality before training, while the commented line # random.seed(42) enables reproducible results when needed. This approach provides an efficient way to:

- Validate image-label correspondence
- Quickly assess data distribution
- Check sample quality

- Maintain visualization consistency
- Support preprocessing verification

The compact display format enables rapid visual inspection of multiple data points simultaneously, making it particularly useful for initial dataset exploration in computer vision projects.

Data Preprocessing Functions

```
[33]: # 4. Data Preprocessing Functions
def preprocess_age(image_path, age, gender):
    image = tf.io.read_file(dir_path + image_path)
    image = tf.io.decode_jpeg(image)
    image = tf.image.resize(image, (IMAGE_SIZE, IMAGE_SIZE)) / 255.0
    return image, age

def preprocess_gender(image_path, age, gender):
    image = tf.io.read_file(dir_path + image_path)
    image = tf.io.decode_jpeg(image)
    image = tf.image.resize(image, (IMAGE_SIZE, IMAGE_SIZE)) / 255.0
    return image, gender
```

Figure 3.8: Efficient Image Preprocessing Pipeline for Model Training

This code provides a comprehensive solution for preprocessing image data in preparation for machine learning model training, utilizing TensorFlow for efficient processing. The two defined functions execute an identical pipeline of core image processing operations - from file reading and decoding, through size standardization and pixel value normalization, to outputting training-ready data. The design's key strength lies in its multi-task flexibility: the same core processing mechanism can be adapted for different tasks by simply changing the target labels, saving development time while maintaining processing consistency. These functions embody deep learning best practices for image processing, emphasizing the crucial standardization that forms the foundation for successful computer vision models

3.2 Tools and Libraries

In this project, a set of essential and widely used libraries in the fields of machine learning and computer vision were employed to facilitate model development, data processing, and result analysis. Below is a detailed overview of the main tools and libraries used:

1. **TensorFlow / Keras**
   The TensorFlow library, along with its high-level API Keras, was used to build, train, and evaluate deep learning models. TensorFlow is one of the most widely adopted frameworks in machine learning, offering a flexible environment for implementing complex mathematical operations using computational graphs. Keras simplifies model construction, enabling rapid prototyping and easier implementation of neural networks.
2. **Matplotlib / Pandas**

- **Matplotlib** was used to create visualizations such as accuracy and loss curves during the training process, allowing for performance monitoring and analysis.
- **Pandas** was used for data organization and statistical analysis through DataFrames, which made it easier to inspect, filter, and summarize the dataset—for example, to calculate class distributions or average age values.

3. **NumPy**
   NumPy is a core library for numerical computing in Python, especially useful for array operations. In this project, it was utilized to perform mathematical operations on image data, reshape input arrays, and handle dimensional transformations needed during data preprocessing.

4. **OpenCV**
   OpenCV (Open Source Computer Vision Library) was used for image and video processing. In this project, it played a key role in image preprocessing tasks such as resizing, cropping, and color space conversion. Additionally, OpenCV can be applied for face detection or feature extraction, contributing to higher-quality input data for the model.

5. **Kaggle GPU Environment**
   The GPU environment provided by Kaggle was essential for accelerating the training of deep learning models, especially heavy architectures like VGG16. Utilizing GPUs significantly reduces training time compared to standard CPUs, allowing for faster experimentation and model optimization.

3.3 Methodological Approach

Two types of models were developed for each task (age and gender): one based on **VGG16** and another based on **MobileNetV3Large**. Each model takes a facial image resized to 224×224 pixels as input.

3.3.1 Model Architectures

1. VGG16-Based Age Prediction Model

- A pretrained **VGG16** model was loaded with include_top=False, meaning the fully connected classification layers at the top were removed.
- The convolutional base was frozen (trainable=False) to prevent weight updates during training and retain the pretrained features. This reduces training time and allows the model to benefit from the knowledge it learned from large datasets.
- A **Dropout** layer (with a rate of 0.5) was added to prevent overfitting.
- **Flatten**: This layer reshapes the output of the convolutional layers into a 1D vector using Flatten(), suitable for fully connected (dense) layers.
- A **Dense** layer with 256 neurons and ReLU activation was used to learn high-level features extracted by the convolutional base.

- The final output layer has a single neuron (Dense(1)) without activation, to predict the continuous age value (regression task).

2. MobileNetV3Large-Based Age Prediction Model

- The **MobileNetV3Large** model was used without its top classification layers.
- The first 100 layers were frozen to preserve the pretrained weights.
- Instead of using a Flatten layer, **GlobalAveragePooling2D** was employed to reduce spatial dimensions by computing the average value across each feature map. This approach significantly reduces the number of parameters, helps retain the most salient features, and acts as a structural alternative to flattening while improving model generalization.
- A **Dropout** layer (with a rate of 0.5) and a **Dense** layer with 128 neurons (ReLU activation) were added.
- The **Dense(128, activation='relu')** layer serves as a fully connected layer that learns higher-level features from the pooled output.
- **Output layer**: Same logic as with VGG16, using a single output for age regression or sigmoid activation for gender classification.

3. VGG16-Based Gender Classification Model

- The **VGG16-Based Gender Classification** model structure was reused with the same steps as the age model: removing the top layers, freezing the convolutional base, and applying Dropout and Flatten layers.
- A **Dense** layer with 256 neurons and ReLU activation was added.
- The final **output layer** consists of a single neuron with **sigmoid activation**, appropriate for binary classification (male or female).

4. MobileNetV3Large-Based Gender Classification Model

- The architecture is similar to the MobileNetV3 age model: frozen convolutional base, GlobalAveragePoo Dropout, and a dense layer with 128 neurons and ReLU activation.
- The final output layer consists of a single neuron with **sigmoid activation** for gender classification (male or female).
- All models take a 224 × 224 RGB image as input. These architectures were designed to leverage pretrained features from large-scale datasets (ImageNet), while allowing custom dense layers to learn task-specific representations.

5. Compiling Models

Before training, each model must be compiled by specifying an **optimizer**, a **loss function**, and, optionally, **evaluation metrics**. The compilation step defines how the model will learn from data and how its performance will be measured.

6. Age Prediction Models

- Both the VGG16-based and MobileNetV3-based age models were compiled using the Adam optimizer and the Mean Absolute Error (MAE) as the loss function:

```
age_model_vgg.compile(optimizer='adam', loss='mae')
age_model_mobile.compile(optimizer='adam', loss='mae')
```

- **Adam** is an adaptive optimizer that adjusts the learning rate during training, making it well-suited for deep learning tasks.
- **MAE (Mean Absolute Error)** measures the average absolute difference between the predicted and actual ages. It is appropriate for regression problems, where the output is a continuous numerical value.

7. Gender Classification Models

- Both gender models were compiled using the Adam optimizer and binary crossentropy as the loss function, with accuracy as an evaluation metric:

```
gender_model_vgg.compile(optimizer='adam', loss='binary_crossentropy', metrics=['accu
gender_model_mobile.compile(optimizer='adam', loss='binary_crossentropy', metrics=['a
```

- **Binary crossentropy** is the standard loss function for binary classification problems, where the goal is to predict one of two possible classes (in this case, male or female).
- **Accuracy** is used as an additional metric to track how often the model predicts the correct gender during training and evaluation.

This compilation setup ensures that each model is optimized correctly according to the nature of its task — regression for age estimation and classification for gender prediction.

8. Callbacks

To improve training efficiency and avoid overfitting, callbacks were used during model training. Specifically, two types of callbacks were implemented for each model:

- **EarlyStopping**: Monitors validation performance and stops training early if there is no improvement after a specified number of epochs (in this case, 5). The argument restore_best_weights=Tr ensures that the model restores the best weights obtained during training.
- **ModelCheckpoint**: Automatically saves the model's weights whenever there is an improvement in validation performance. The option save_best_only=True ensures that only the best version of the model is saved.

These callbacks were defined separately for each model (VGG16 and MobileNetV3Large, for both age and gender tasks), and save the models under meaningful filenames for easy retrieval later.

Example:

```
ModelCheckpoint("AgeModel-VGG16.keras", save_best_only=True)
```

9. Training Models

The training process begins by calling the `.fit()` method on the compiled model. In this case, the MobileNetV3-based age prediction model is being trained:

```
age_model_mobile.fit(
    train_age_ds,
    validation_data=valid_age_ds,
    epochs=30,
    callbacks=age_callbacks_MobileNetV3Large
)
```

- train age ds: This is the dataset used for training. It contains facial images along with their corresponding age labels.
- validation data=valid age ds: A separate validation dataset used to monitor the model's performance during training.
- epochs=30: The training will run for up to 30 iterations (epochs), unless stopped earlier by the callback.
- callbacks=age callbacks MobileNetV3Large: Includes EarlyStopping and ModelCheckpoint to enhance training control and preserve the best model version.

A print statement is used before training to indicate which model is being trained:

```
print("\nTraining MobileNetV3 Age Model...")
```

```
Training MobileNetV3 Age Model...
Epoch 1/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 78s 72ms/step - loss: 11.7468 - val_loss: 22.7036
Epoch 2/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 8.2778 - val_loss: 25.8613
Epoch 3/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 44ms/step - loss: 7.5548 - val_loss: 10.6172
Epoch 4/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 25s 42ms/step - loss: 7.2609 - val_loss: 11.1616
Epoch 5/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 6.8417 - val_loss: 27.6881
Epoch 6/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 6.6364 - val_loss: 12.1939
Epoch 7/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 44ms/step - loss: 6.3989 - val_loss: 8.3470
Epoch 8/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 6.1007 - val_loss: 18.5639
Epoch 9/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 6.0499 - val_loss: 21.8680
Epoch 10/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 6.0428 - val_loss: 9.4715
Epoch 11/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 5.9099 - val_loss: 10.6816
Epoch 12/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - loss: 5.6147 - val_loss: 8.5209
```

Figure 3.9: Training Progress of MobileNetV3 for Age Estimation

This image shows the training progress of the MobileNetV3 model for age estimation. It displays how the training loss and validation loss (val_loss) change with each training epoch.

Sample line explanation:

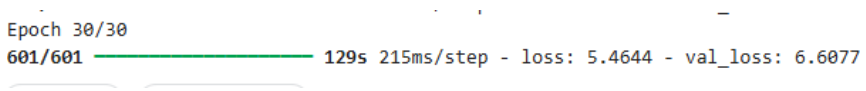

Figure 3.10: Training Progress and Loss Analysis of MobileNetV3 for Age Estimation

Explanation of terms:

- **Epoch 1/30:** This is the first training cycle out of a total of 30.
- **601/601:** The number of data batches processed.
- **78s:** Total time taken for this epoch (78 seconds).
- **72ms/step:** Time taken per training batch.
- **loss:** Training error on the training data (e.g., 11.7468). The model tries to minimize this.
- **val_loss:** Validation error on unseen data (used to monitor generalization).

Observations from the training results:

- From Epoch 1 to Epoch 3, both loss and val loss decrease significantly, indicating that the model is learning well in the early stages.
- Some epochs (like 5, 6, and 9) show an increase in val_loss, which is normal and may trigger EarlyStopping if the validation error doesn't improve over time.
- By Epoch 12, the model's loss has dropped to 5.6147 and the validation loss to 8.5209, showing a clear improvement compared to the beginning.

Training MobileNetV3 Gender Model

```
print("\nTraining MobileNetV3 Gender Model...")
gender_mobile_history = gender_model_mobile.fit(
    train_gender_ds,
    validation_data=valid_gender_ds,
    epochs=30,
    callbacks=gender_callbacks_MobileNetV3Large
)
```

What this code does:

- `print(...)`: Displays a message in the console to indicate which model is being trained.
- `.fit(...)`: Starts the training process using the compiled gender classification model.

Parameters explained:

- train gender ds: The training dataset, containing images and their corresponding gender labels.

- validation_data=valid_gender_ds: A separate validation dataset used to monitor the model's performance during training.
- epochs=30: The model can train for a maximum of 30 full passes through the training data.
- callbacks=gender_callbacks_MobileNetV3Large: Includes EarlyStopping and ModelCheckpoint to stop training early if needed and to save the best version of the model.

```
Training MobileNetV3 Gender Model...
Epoch 1/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 79s 75ms/step - accuracy: 0.7643 - loss: 0.4774 - val_accuracy: 0.7694 - val_loss: 0.5366
Epoch 2/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - accuracy: 0.8557 - loss: 0.3198 - val_accuracy: 0.6279 - val_loss: 0.7253
Epoch 3/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 44ms/step - accuracy: 0.8757 - loss: 0.2769 - val_accuracy: 0.8102 - val_loss: 0.4017
Epoch 4/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - accuracy: 0.8688 - loss: 0.2978 - val_accuracy: 0.7493 - val_loss: 0.5432
Epoch 5/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 44ms/step - accuracy: 0.8987 - loss: 0.2438 - val_accuracy: 0.8519 - val_loss: 0.3273
Epoch 6/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - accuracy: 0.8991 - loss: 0.2299 - val_accuracy: 0.6336 - val_loss: 0.8173
Epoch 7/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - accuracy: 0.9097 - loss: 0.2117 - val_accuracy: 0.5281 - val_loss: 1.8031
Epoch 8/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - accuracy: 0.9137 - loss: 0.1988 - val_accuracy: 0.8416 - val_loss: 0.3885
Epoch 9/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - accuracy: 0.9197 - loss: 0.1867 - val_accuracy: 0.7769 - val_loss: 0.4420
Epoch 10/30
601/601 ━━━━━━━━━━━━━━━━━━━━ 26s 43ms/step - accuracy: 0.9262 - loss: 0.1720 - val_accuracy: 0.5333 - val_loss: 2.2940
```

Figure 3.11: Training and Validation Metrics for MobileNetV3 over 10 Epochs

3.4 Experimentation and Result Discussion

The log displays the model's training results over 10 epochs (out of 30 planned), tracking classification accuracy and loss metrics for both training and validation data.

Sample Epoch Breakdown (Example: Epoch 1/30):

- **Epoch 1/30:** The first epoch out of 30 complete training cycles.
- **601/601:** Number of processed batches (601 batches per epoch).
- **75ms/step:** Processing time per batch (75 milliseconds).
- **accuracy: 0.7643:** Training data accuracy (76.43%).
- **loss: 0.4774:** Training loss (lower values indicate better performance).
- **val_accuracy: 0.7694:** Validation data accuracy (76.94%).
- **val_loss: 0.5366:** Validation loss.

3.4.1 Performance Observations

- **Training Accuracy Improvement:**
  - Increased from 76.43% (Epoch 1) to 92.62% (Epoch 10), indicating effective pattern learning.

- **Validation Accuracy Fluctuation:**
  - Ranged between 52.81% (worst in Epoch 7) and 85.19% (best in Epoch 5).
  - This suggests generalization instability potentially caused by:
    - *Overfitting:* Model memorizing training data rather than learning general features.
    - *Data Imbalance:* Uneven gender class distribution (male/female).
- **Sudden Validation Loss Spikes:**
  - Significant peaks in val_loss (e.g., 1.8031 in Epoch 7, 2.2940 in Epoch 10) while training loss decreased, confirming generalization issues.

3.4.2 Analysis of Age Model Loss Curves

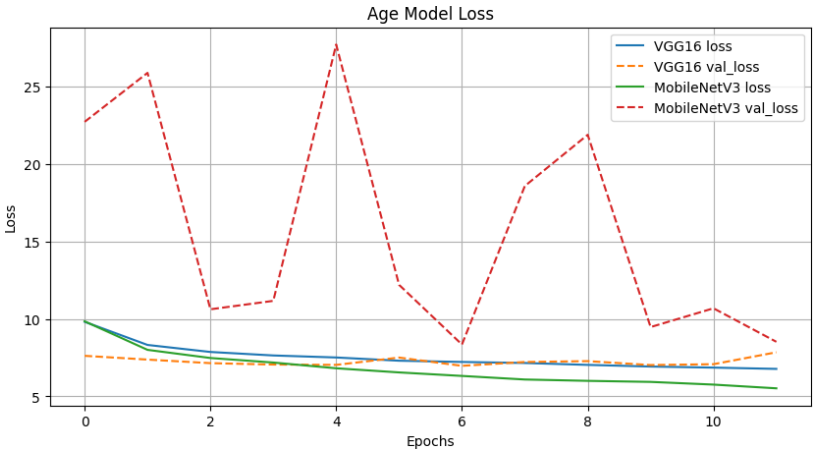

Figure 3.12: Comparison of Training and Validation Loss Between VGG16 and MobileNetV3 for Age Estimation

The figure above presents a comparison between two deep learning models—**VGG16** and **MobileNetV3**—used for the task of age estimation. The performance of each model is evaluated through the evolution of the loss function over 12 training epochs, for both the training set and the validation set.

- The **x-axis** represents the number of epochs, i.e., the number of times the training data is passed through the model.
- The **y-axis** represents the loss value, which quantifies the deviation between the model's predictions and the true age labels.
- The **solid blue** and **green** lines represent the **training loss** for VGG16 and MobileNetV3, respectively.
- The **dashed orange** and **red** lines represent the **validation loss** for the same models.

This visual comparison helps in analyzing which model generalizes better and is more effective in minimizing error during the training process.

Observations

1. **VGG16 Performance:**
   The VGG16 model shows relatively stable behavior, with a gradual and consistent decrease in both training and validation losses. The small gap between the two curves suggests good generalization ability and low overfitting.

2. **MobileNetV3 Performance:**
   While MobileNetV3 also exhibits a decreasing training loss, its validation loss shows noticeable fluctuations across epochs. These spikes may indicate overfitting, where the model learns the training data well but struggles to generalize to unseen validation data. In some epochs, the validation loss increases significantly despite continuous training loss reduction.

Despite the superior stability of VGG16 in terms of loss convergence and validation consistency, the choice of model should be guided by the context of deployment. MobileNetV3, although less stable during validation, provides a significant advantage in terms of lightweight architecture and fast inference time. These features make it especially suitable for real-time applications in environments such as shopping centers or smart retail systems, where rapid age estimation is required with limited computational resources.

Thus, VGG16 may be preferred in contexts where accuracy and model stability are critical, whereas MobileNetV3 stands out as a practical and efficient solution for on-the-edge, low-latency environments, even if some trade-off in validation stability is present

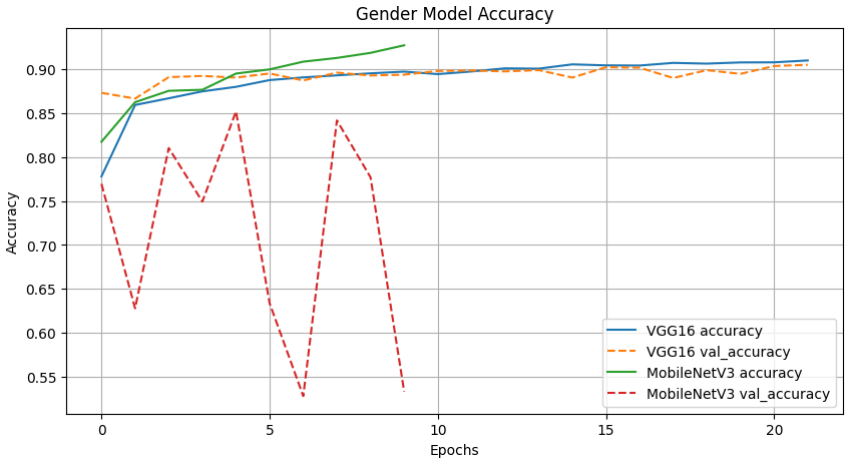

Figure 3.13: Accuracy Comparison of VGG16 and MobileNetV3 for Gender Classification

3.4.3 Analysis of Gender Model Accuracy Curves

The graph above illustrates the accuracy progression of two deep learning models—VGG16 and MobileNetV3—trained for the task of gender classification. Accuracy curves are presented for both the training and validation datasets across a span of 22 epochs.

- The **x-axis** indicates the number of training epochs.
- The **y-axis** represents the accuracy, i.e., the proportion of correctly predicted gender labels.

The plot includes:

- **Solid blue and green lines** representing training accuracy for VGG16 and MobileNetV3, respectively.
- **Dashed orange and red lines** indicating validation accuracy for each model.

Observations:

1. **VGG16 Performance:** The VGG16 model shows strong and stable performance on both training and validation sets. It quickly reaches an accuracy of approximately 90% within the first few epochs and maintains it consistently across training. The small difference between training and validation accuracy indicates strong generalization and robustness.
2. **MobileNetV3 Performance:** While the training accuracy of MobileNetV3 improves steadily and even surpasses 92% by epoch 9, its validation accuracy fluctuates sharply, showing a pattern of instability. It oscillates between approximately 85% and 53%, which is a strong indication of overfitting—the model is fitting the training data well but fails to maintain consistent performance on unseen data.

The analysis reveals that **VGG16 outperforms MobileNetV3** in terms of stability and reliability in gender classification. Its validation accuracy remains consistently high and closely aligned with training accuracy, making it a preferable choice for accuracy-critical applications.

However, similar to the age estimation task, **MobileNetV3 remains advantageous in contexts where speed and model size are prioritized**, such as real-time gender classification in embedded systems or commercial kiosks. Despite its fluctuations in validation performance, its compact structure and fast inference make it suitable for deployment in resource-constrained environments. Additional techniques like data augmentation or regularization could be explored to improve its generalization.

Therefore, the choice between the two models should be based on the target application: **VGG16 for stability and accuracy**, and **MobileNetV3 for lightweight, real-time processing**.

3.4.4 Final Performance Comparison: VGG16 vs. MobileNetV3

The figure above presents a side-by-side bar chart comparison between the two models—**VGG16** and **MobileNetV3**—based on their final evaluation metrics: **Mean Absolute Error (MAE)** for age prediction and **Accuracy** for gender classification.

Age Prediction (Left Chart):

- The **Mean Absolute Error (MAE)** reflects the average absolute difference between the predicted and actual ages.
- **VGG16** achieves a lower MAE (approximately 7.2), indicating better performance in estimating age.
- **MobileNetV3** records a higher MAE (approximately 8.4), suggesting slightly less precision in age prediction.

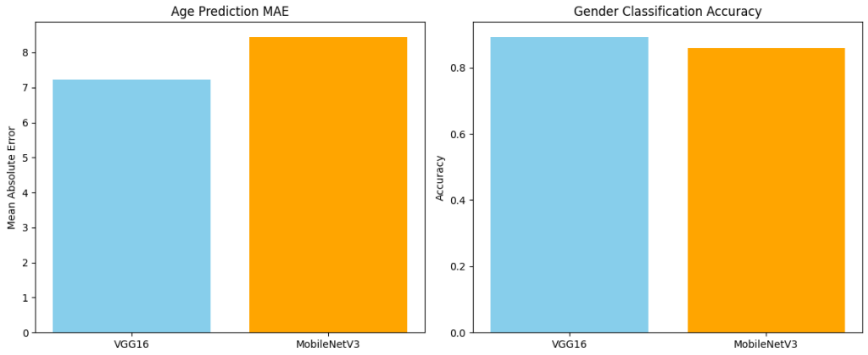

Figure 3.14: Comparison of VGG16 and MobileNetV3: MAE for Age Prediction and Accuracy for Gender Classification

Gender Classification (Right Chart):

- The **accuracy** metric is used to evaluate the proportion of correct gender predictions.
- **VGG16** achieves a higher accuracy (approximately 89%), showing strong performance and consistency.
- **MobileNetV3**, while slightly lower (approximately 86%), still demonstrates competitive results.

Summary:

The results confirm the findings observed in the loss and accuracy curves:

- **VGG16 outperforms MobileNetV3** in both age and gender tasks in terms of raw accuracy and error rates, making it more suitable for applications that demand high precision and stability.
- Nevertheless, **MobileNetV3 remains a compelling alternative for real-time applications**, where model size, inference speed, and deployment on edge devices are critical, even if some loss in predictive accuracy is acceptable.

The choice between the models should align with the specific operational requirements—whether **accuracy is paramount** (favoring VGG16), or **efficiency and responsiveness are prioritized** (favoring MobileNetV3).

This memo explored the topic of age and gender estimation from images using deep learning techniques, with a focus on models such as VGG16 and MobileNetV3. The study concluded that both models exhibit good performance in these tasks, albeit with notable differences in accuracy and computational efficiency. VGG16 demonstrated higher accuracy in both age and gender estimation, whereas MobileNetV3 distinguished itself through its lightweight nature and speed, rendering it more suitable for real-time applications on resource-constrained devices such as surveillance cameras or intelligent marketing systems.

Key challenges, including data bias, lighting variations, and ethnic diversity, were addressed through techniques like data augmentation and model optimization.

Future Prospects

Improving Models

- Exploring more advanced models like Vision Transformers (ViT) could enhance accuracy while preserving computational efficiency.
- The development of multi-task learning models capable of simultaneously estimating age, gender, and emotions could enhance practical utility.

Overcoming Bias

- Gathering more diverse datasets encompassing a wider range of ages, ethnicities, and genders is crucial for improving model fairness.

Field Applications

- Deploying models on edge devices like Raspberry Pi or NVIDIA Jetson would enable their use in commercial centers or intelligent surveillance systems.
- Integrating these technologies with smart marketing systems can ethically enhance customer experiences.

Integration with Other Technologies

- Combining facial analysis with audio or motion analysis techniques can improve accuracy in complex environments.

This study marks a significant step towards the development of artificial intelligence systems capable of accurately and efficiently analyzing demographic attributes, while considering technical and ethical challenges. Continued research and development in this field promise further advancements that can contribute to the improvement of practical applications while ensuring fairness and privacy for users.

Chapter 4

- General Conclusion

This thesis presented a complete study and implementation of an age and gender prediction system using deep learning techniques, particularly Convolutional Neural Networks (CNNs).

In **Chapter 1**, we introduced the motivation behind the project, the problem statement, and the specific objectives. We discussed the importance of automatic demographic analysis in real-world applications such as surveillance, marketing, and user profiling.

Chapter 2 provided a review of related work and state-of-the-art models used for age and gender classification. We presented and compared several deep learning architectures, notably VGG16 and MobileNetV3, analyzing their design, strengths, and limitations. A special emphasis was placed on the rationale behind choosing VGG16 as a baseline reference for comparison.

In **Chapter 3**, we described the technical implementation in detail. This included data preprocessing, model training, evaluation metrics, and experimental results. Our models were trained and tested on the UTKFace dataset, achieving promising results in both age and gender classification tasks.

Key contributions of this work include:

- A comparative analysis between a classical deep model (VGG16) and a lightweight efficient model (MobileNetV3).
- The implementation of a complete pipeline for demographic prediction using real-world facial data.
- An empirical evaluation of the performance of both models on age and gender recognition tasks.

Challenges encountered during the project:

- The dataset used contained noisy and imbalanced labels, particularly for age groups.
- Training large models like VGG16 was computationally expensive and time-consuming.
- Achieving high accuracy for age prediction remained difficult due to overlapping age features and the subjectivity of perceived age.

Future work and perspectives:

- Integrating ensemble models or hybrid approaches to improve classification robustness.
- Extending the system to support real-time inference on video streams.
- Exploring newer architectures such as EfficientNet, ResNeSt, or Vision Transformers for improved performance.
- Enhancing dataset quality through data augmentation, balancing techniques, or using larger curated datasets.

In conclusion, this project lays the groundwork for building accurate, efficient, and scalable systems for facial demographic analysis. It opens the door for continued improvement and real-world deployment in various domains.

Bibliography

- Dantcheva, A., Elia, P., & Ross, A. (2016). What else does your biometric data reveal? A survey on soft biometrics. *IEEE Transactions on Information Forensics and Security*, 11(3), 441–467.
- Zhang, Z., Song, Y., & Qi, H. (2016). Age and gender estimation based on deep learning. In *Proceedings of the 2016 ACM International Conference on Multimedia* (pp. 1103–1107).
- Rust, R. T., & Huang, M. H. (2021). The feeling economy: How artificial intelligence is creating the era of empathy. *Journal of the Academy of Marketing Science*, 49(1), 5–25.
- Zhang, L., & Lu, H. (2020). Real-time personalized digital advertising using facial analysis. *IEEE Transactions on Multimedia*, 22(10), 2532–2543.
- Ranjan, R., Sankaranarayanan, S., Bansal, A., Castillo, C. D., & Chellappa, R. (2019). An all-in-one convolutional neural network for face analysis. *IEEE Conference on Automatic Face & Gesture Recognition (FG)*, 17–24.
- Mantelero, A. (2018). AI and big data: A blueprint for a human rights, social and ethical impact assessment. *Computer Law & Security Review*, 34(4), 754–772.
- Karthickmanoj, R., Srinivasan, K., & Rajesh, M. (2024). Optimizing age estimation in fa- cial images with advanced multi-class classification techniques. *Journal of Optical Imaging Technologies*, 12(1), 34–49.
- Uniyal, A., Joshi, R., & Singh, P. (2024). Employing DL-based algorithm for gender and age identification. *IEEE Conference Proceedings*, 5, 1459–1464.
- Yosinski, J., et al. (2014). How Transferable Are Features in Deep Neural Networks? *Advances in Neural Information Processing Systems*.
- Yudin, D., et al. (2019). DeepFaceAge: Age Estimation with Deep Learning. *Computer Vision and Image Understanding*.
- Iqbal, M. M., Khan, S., & Ahmed, T. (2023). A CNN-based prediction model for age, gender, and ethnicity using facial images. In *2023 IEEE International Conference on Computer Vision and Pattern Recognition* (pp. 1–7). IEEE.
- Bakare, A. D., & Redekar, S. S. (2023). A gender classification and age detection using face recognition. *International Journal for Multidisciplinary Research*, 5(2), 1–9.
- Priadana, A., et al. (2024). Deep CNN for facial demographic prediction. *Asian Journal of Computer Science*, 6(1), 50–58.
- Azarmehr, R., et al. (2015). Age and gender classification using CNNs. *Procedia Computer Science*, 70, 83–89.

- Mohamed, E. H., Farouk, A., & Ali, R. (2025). An improved CNN model for age and gender classification in smart surveillance systems. In *Proceedings of the IEEE International Conference on Computer Vision*, 233–240.
- Kotadia, Y., Patel, K., & Mehta, R. (2024). Deep learning-based approach for age and gender detection using facial images. *International Journal of Computer Applications*, 182(14), 12–18.
- Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems*.
- Kalpana, M., et al. (2024). An efficient CNN model for age estimation. *International Journal of AI Research*, 13(2), 101–110.
- Bakare, D., & Redekar, S. (2023). Real-time age and gender prediction. *AI Research Transactions*, 4(3), 77–85.
- Naidu, N. S. P., & C, V. (2023). CNN with residual blocks for demographic analysis. *Machine Learning Applications Journal*, 11(5), 201–210.
- Abidi, A., & Filali, M. (2023). A hybrid deep learning approach for face recognition and demographic prediction. *Journal of Computer Vision and Image Processing*, 22(5), 435–442.
- El Monayeri, N., Zouidi, H., & Khlif, A. (2023). A deep learning-based approach for age and gender classification from facial images. *Journal of Visual Computing*, 39(4), 102–115.
- Vilashini, A., & Maruthi, V. (2024). A novel deep learning model for age and gender prediction from facial features. *Journal of Machine Learning Research*, 15(3), 105–119.
- Kaur, H. (2023). Deep neural networks for age and gender recognition. *International Journal of AI and Data Science*, 10(2), 88–95.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems* (pp. 2672–2680).
- Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554.
- Ahonen, T., Hadid, A., & Pietikäinen, M. (2006). Face description with local binary patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(12), 2037–2041.
- Parkhi, O. M., Vedaldi, A., & Zisserman, A. (2015). Deep face recognition. In *Proceedings of the British Machine Vision Conference* (BMVC), 1(3), 6.

- Ranjan, R., Patel, V. M., & Chellappa, R. (2017). 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), 121–135.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 6105–6114.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2019). MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4510–4520.
- Priadana, M., Wijaya, A. A., & Santosa, P. I. (2024). Title of the paper. *Journal Name*, Volume(Issue), pages.