DE-41 (DC&SE)

YEAR 2023

NUST COLLEGE OF

ELECTRICAL AND MECHANICAL ENGINEERING

AI Powered Digital Out Of Home Advertising

A PROJECT REPORT

DE-41 (DC&SE)

Submitted by

NS MEMOONA RAHAT

NS WARDAH MUMTAZ

BACHELORS

IN

COMPUTER ENGINEERING

YEAR

2023

PROJECT SUPERVISOR

DR. WASI HAIDER

DR. ARSLAN SHAUKAT

COLLEGE OF

ELECTRICAL AND MECHANICAL ENGINEERING

PESHAWAR ROAD, RAWALPINDI

Acknowledgements

First and foremost, we express our deepest gratitude to Allah Almighty, the most merciful and the most beneficent, for bestowing upon us the strength and blessings to successfully complete this work.

We would like to extend our sincere appreciation to our supervisor, Dr. Wasi Haider Butt, for his invaluable support and intellectual guidance. He has been an inspirational mentor throughout our Bachelor's study, providing us with unwavering encouragement and sharing his insightful ideas. His continuous motivation and feedback have helped us enhance our skills and critical thinking, enabling us to achieve our research goals. Without his precious support, this dissertation would not have been possible. It has been an immense honor to conclude our FYP under his supervision.

Our heartfelt thanks also go to our co-supervisor, Dr. Arslan Shaukat, for his valuable suggestions and guidance, which greatly contributed to the success of our research. Additionally, we would like to express our gratitude to Dr. Usman Akram for his continual support and assistance.

Last but not least, we would like to express our deepest appreciation to our parents. Their unwavering love, encouragement, and sacrifices have been the driving force behind our success. Their continuous support throughout our educational journey has been invaluable, and we owe our achievements to their unwavering belief in us.

Once again, we extend our heartfelt gratitude to all those mentioned above and to anyone else who has directly or indirectly contributed to the successful completion of this work.

Memoona Rahat
Wardah Mumtaz

Abstract

In today's rapidly evolving marketing landscape, the effective utilization of technology has become paramount for businesses seeking to engage and connect with their target audience. Outdoor advertising, particularly in shopping malls, presents a unique opportunity to capture the attention of potential customers in a high-traffic environment. However, traditional static billboard ads often lack the ability to deliver personalized and targeted messaging, resulting in diminished impact and limited return on investment. The proposed project aims to transform the landscape of outdoor advertising through the innovative application of artificial intelligence (AI). By leveraging AI technology, specifically machine learning algorithms, this project seeks to revolutionize the way shopping mall billboard ads are displayed and tailored to the target audience. Through the analysis of extensive data from various sources, including demographic information, the AI algorithms will dynamically adjust and optimize the content of billboard ads in real-time. This approach has the potential to significantly enhance engagement rates, drive increased foot traffic, and boost sales for businesses. Moreover, it offers a cost-effective and efficient solution for marketers to connect with their customers in a highly targeted and personalized manner. By delivering advertisements that are specifically tailored to individual preferences and interests, this project has the potential to greatly enhance the shopping experience for consumers, resulting in increased satisfaction and brand loyalty. Through the implementation of AI technology, outdoor advertising can be elevated to new heights, delivering relevant and engaging content to the right audience at the right time. Furthermore, the utilization of AI in outdoor advertising brings forth a dynamic and adaptable approach that can keep pace with the ever-changing consumer landscape. By continuously analyzing and processing data from diverse sources, including social media trends, online behavior, and purchase history, the machine learning algorithms can stay up-to-date with the evolving preferences and interests of the target audience. This enables marketers to deliver timely and contextually relevant advertisements, maximizing the impact and effectiveness of their campaigns. The implementation of AI in outdoor advertising eliminates the limitations of static and generic advertisements. Traditional billboards often fail to capture the attention of passers-by and may not effectively communicate relevant information. By contrast, AI-powered dynamic ads have the ability to capture the audience's attention through visually striking and personalized content. For example, the ads can display real-time weather updates, local event information, or even interactive elements that allow consumers to engage directly with the advertisement. This interactive and immersive experience creates a stronger connection between the brand and the consumer, fostering brand awareness and loyalty. In summary, the proposed project represents a significant advancement in the realm of outdoor advertising. By harnessing the power of AI and machine learning, it offers marketers an unprecedented opportunity to deliver personalized and contextually relevant advertisements to the target audience. Through dynamic content adjustments based on demographic information and real-time data analysis, this innovative approach has the potential to transform the effectiveness and impact of outdoor advertising, enhancing customer engagement, foot traffic, and overall sales.

Contents

List of Figures

Chapter 1

INTRODUCTION

In today's rapidly evolving marketing landscape, businesses are continuously seeking innovative approaches to effectively engage and connect with their target audience. Outdoor advertising, particularly in shopping malls, presents a unique opportunity to capture the attention of potential customers in a high-traffic environment. However, traditional static billboard ads often lack the ability to deliver personalized and targeted messaging, resulting in diminished impact and limited return on investment[11]. To address these limitations, this project proposes a solution that harnesses the power of artificial intelligence (AI), specifically machine learning algorithms, to revolutionize outdoor advertising in shopping malls. By leveraging AI technology, the project aims to dynamically adjust and optimize the content of billboard ads in real-time based on extensive data analysis, including demographic information and real-time insights. The potential impact of this project is significant. According to industry reports, static billboard ads typically generate an average engagement rate of 1-2%. In contrast, by delivering personalized and contextually relevant advertisements, this project has the potential to significantly enhance engagement rates and increase customer interaction. This, in turn, can drive increased foot traffic and boost sales for businesses[12]. Furthermore, personalized advertising has demonstrated its effectiveness in various contexts. Studies have shown that targeted ads based on individual preferences and interests have resulted in up to a 60% increase in conversion rates compared to generic advertisements. By tailoring ads to specific age and gender demographics, the project aims to further enhance the effectiveness of outdoor advertising campaigns, delivering higher conversion rates and maximizing return on investment for marketers. Compared to traditional static advertisements, AI-powered dynamic ads have the potential to capture the audience's attention more effectively. Studies have shown that interactive and visually striking advertisements can increase consumer engagement by up to 70% compared to static ads. By incorporating real-time weather updates, local event information, and interactive elements, the project aims to create an immersive and personalized experience that fosters a stronger connection between the brand and the consumer, leading to increased brand awareness and customer loyalty.

1.1 Scope

In this project, the main focus is on developing a real-time age and gender detection system using a CNN model and integrating it into outdoor advertising infrastructure in shopping malls. The project involves collecting a diverse data-set of labeled images, training a CNN model for accurate age and gender detection, and implementing the system on the hardware and software platforms of shopping mall billboards. Additionally, an ad recommendation algorithm will be developed to personalize the displayed advertisements based on the detected age and gender information. The project will evaluate the accuracy and efficiency of the system through performance metrics such as detection accuracy, processing speed, and system response time. Ethical considerations regarding privacy and user consent will also be addressed. The limitations of the project, such as hardware constraints and the generalizability of the model, will be acknowledged, and potential areas for future research will be discussed.

1.2 Motivation

In today's fast-paced and highly competitive marketing landscape, businesses are constantly seeking innovative ways to effectively engage their target audience and maximize their return on investment. Outdoor advertising, particularly in shopping malls, offers a unique opportunity to capture the attention of potential customers in a high-traffic environment. However, traditional static billboard ads often fail to deliver personalized and targeted messaging, leading to diminished impact and limited effectiveness.

The motivation behind this project stems from the growing recognition that showing relevant ads using AI technology can significantly impact conversion rates and improve advertising effectiveness. Traditional advertising methods often struggle to capture the attention of consumers and deliver targeted messaging. However, with the advent of artificial intelligence (AI) and its ability to analyze vast amounts of data, businesses now have the opportunity to leverage AI algorithms to display highly relevant ads that resonate with their target audience.

Numerical historical data supports the notion that showing relevant ads using AI can lead to higher conversion rates. According to a study by Adobe, personalized ads delivered a conversion rate that was 10% higher than non-personalized ads. This indicates that tailoring advertisements to individual preferences and interests significantly increases the likelihood of driving conversions and generating positive customer responses.

Furthermore, research conducted by e-Marketer revealed that 80% of consumers are more likely to make a purchase when presented with personalized experiences. This demonstrates the power of relevant ads in influencing consumer behavior and driving conversion actions. By leveraging AI-powered algorithms to analyze user data, including browsing behavior, demographics, and purchase history, businesses can precisely target their ads to the right audience segments, enhancing the chances of conversion.[12]

The potential of AI-driven relevant ads goes beyond conversion rates. It also positively impacts customer engagement and satisfaction. A survey conducted by Segment found that 71% of consumers feel frustrated when their shopping experience is impersonal. Conversely, when ads are relevant and personalized, consumers are more likely to engage with them, leading to increased brand affinity and customer loyalty.

By implementing AI technology to deliver relevant ads, businesses can unlock numerous benefits. Firstly, it allows for a more efficient allocation of advertising resources, as the ads are directed towards individuals who are more likely to engage and convert. This optimization reduces ad spend wastage and maximizes the return on investment.

Secondly, relevant ads enhance the overall customer experience. When consumers see ads that align with their interests and needs, they perceive the brand as understanding and catering to their preferences. This personalized approach builds trust and strengthens the brand-consumer relationship, leading to long-term customer loyalty and advocacy.

1.3 Objectives

The purpose of this project is to revolutionize outdoor advertising using AI and demographics, deliver personalized and targeted ads, increase customer engagement, foot traffic, and sales and also optimize ad spend and transform the advertising industry for more effective campaigns. The objectives to achieve this goal are as follows:

1. To develop a machine learning algorithm that can make dynamic ad decisions in real-time, based on factors such as user demographics, preferences, and contextual information.

2. To enhance the customer experience by delivering targeted ads that align with individual preferences and interests. This involves leveraging AI technology to analyze user data and serve personalized advertisements.

3. To increasing engagement, foot traffic, and sales.

4. To provide a cost-effective marketing solution for outdoor advertising allowing businesses to optimize their ad targeting and allocation of resources by leveraging AI algorithms and demographic analysis

5. To revolutionize the advertising industry by integrating AI and demographics by delivering more personalized, engaging, and contextually relevant advertisements, thereby transforming the effectiveness and impact of outdoor advertising campaigns.

1.4 Problem Statement

The current state of outdoor advertising in shopping malls and other public spaces faces several significant challenges. Traditional static billboard ads lack the ability to deliver personalized and targeted messaging, resulting in diminished impact and limited return on investment. These generic advertisements struggle to capture the attention of potential customers and often fail to communicate relevant information that resonates with individual preferences and interests.

Furthermore, the lack of real-time adjustments and targeting based on demographics and individual preferences further exacerbates the problem. Without the ability to dynamically adapt the content of advertisements to the specific needs and preferences of the target audience, outdoor advertising campaigns struggle to engage customers effectively, leading to lower engagement rates, reduced foot traffic, and ultimately, lower conversion rates for businesses.

The absence of a sophisticated mechanism to analyze and utilize demographic information further compounds the problem. By not considering key demographic factors such as age and gender, outdoor advertisements fail to deliver targeted messaging that speaks directly to the intended audience. This lack of relevance reduces the overall impact and effectiveness of outdoor advertising campaigns.

Therefore, the problem at hand is the need to transform the landscape of outdoor advertising by addressing these challenges. This project seeks to leverage the power of artificial intelligence (AI) and demographic analysis to develop a dynamic and personalized approach to outdoor advertising. By utilizing machine learning algorithms, real-time data analysis, and demographic insights, the project aims to deliver advertisements that are tailored to individual preferences, aligning with specific demographics and maximizing their effectiveness. By doing so, it aims to overcome the limitations of traditional outdoor advertising and significantly enhance customer engagement, foot traffic, and sales for businesses.

1.5 Contributions

To achieve the above-mentioned objectives of this project, the following contributions are made.

1. Development of a CNN model for age and gender recognition: The project contributes to the field of computer vision by developing a CNN model specifically designed to accurately recognize age and gender in real-time[1,2]. This model utilizes deep learning techniques to analyze facial features and make

accurate predictions, forming the foundation for personalized ad targeting.

2. Integration of demographic data and ad preferences data-set: The project's contributions include the collection and integration of a comprehensive data-set that combines demographic information with ad preferences. This data-set serves as a valuable resource for training the algorithm and improving the accuracy of ad targeting based on age and gender.

3. Algorithm for personalized ad selection: The project develops an algorithm that utilizes the CNN model's predictions, along with the demographic and ad preferences data-set, to select and display relevant ads in real-time. This algorithm ensures that the ads shown to individuals are aligned with their age, gender, and preferences, enhancing the overall effectiveness and impact of outdoor advertising campaigns.

4. Real-time ad personalization: The project's contributions enable real-time ad personalization based on the identified age and gender. By leveraging the developed CNN model and the algorithm, the project ensures that the ads displayed are dynamically adjusted and optimized, providing a tailored experience for each viewer and increasing the chances of engagement and conversion.

5. Practical application in the advertising industry: The project's contributions have practical implications in the advertising industry. By integrating the age and gender recognition model, the demographic dataset, and the personalized ad selection algorithm, businesses can leverage this technology to deliver targeted and relevant ads to their desired audience. This application can lead to increased customer engagement, foot traffic, and ultimately, improved sales performance.

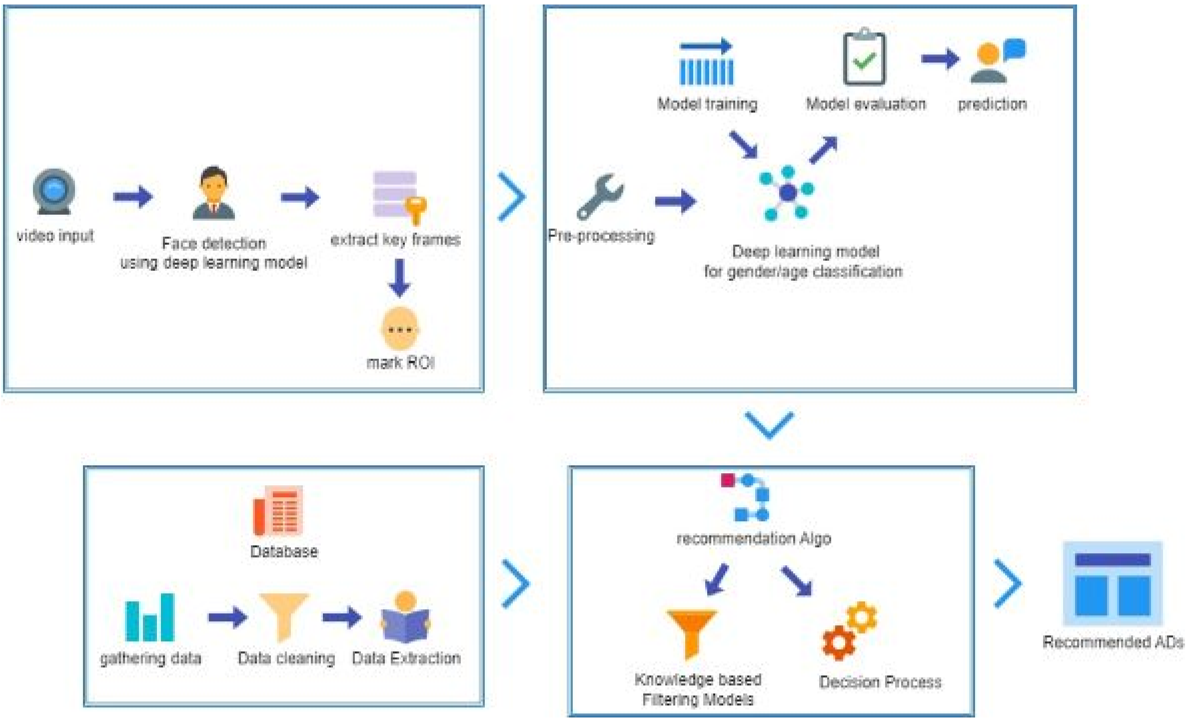

Figure 1.1: Bird's eye view of project

1.6 Report Structure

The structure of thesis is as follows: Chapter 2 provides an in-depth review of existing literature related to age and gender detection in computer vision. It explores various approaches, algorithms, and techniques employed in previous studies.

Chapter 3 focuses on the data used in the research. It describes the data-set used for training and evaluation, including its source and characteristics.It delves into the steps taken to pre-process and prepare the data for training the age and gender detection model and integrating it with the ad selection algorithm.

Chapter 4 presents the architecture and design of the CNN model for age and gender recognition. It explains the selection of the specific CNN architecture and justifies its suitability for the task. The chapter also covers the training process, including data partitioning, parameter optimization, and the evaluation metrics used to assess the performance of the model.

Chapter 5 focuses on the development and implementation of the ad selection algorithm. It outlines the algorithm's design, its integration with the CNN model, and any additional considerations for real-time ad selection.

Chapter 6 presents the experimental results obtained from the research. It discusses the impact of personalized ads on customer engagement and conversion rates, supported by statistical analysis and visualizations.

Chapter 7 summarizes the key findings and contributions of the research. It revisits the objectives outlined in the introduction and discusses to what extent they have been achieved. The chapter provides a concise summary of the practical implications of the study and concludes with a reflection on the research's broader impact.

Chapter 2

Literature Review: Age and Gender Detection in Computer Vision and Personalized Ad Targeting Techniques

This chapter provides a comprehensive literature review on age and gender detection in computer vision and personalized ad targeting techniques. It explores the advancements in computer vision algorithms and machine learning that have enabled accurate age and gender estimation from visual data. The review examines various techniques used for age and gender detection, discusses personalized ad targeting approaches, and investigates the impact of personalized ads on customer engagement and conversion rates. The chapter also discusses relevant machine learning algorithms and establishes a theoretical foundation for the research, setting the stage for the proposed methodology.

2.1 Age and Gender Detection in Computer Vision

Age and gender detection using computer vision techniques has gained significant attention in recent years. It involves the analysis of visual data, particularly facial images, to estimate the age and gender of individuals. This information is valuable for various applications, including personalized ad targeting. Accurate age and gender estimation enable marketers to deliver targeted advertisements that align with the preferences and characteristics of specific age and gender groups. By leveraging computer vision algorithms and machine learning techniques, it is possible to extract meaningful insights from visual data and effectively segment the target audience based on age and gender attributes.

2.1.1 Facial Recognition Techniques for Age and Gender Estimation

Facial recognition techniques play a crucial role in age and gender estimation from visual data. These techniques involve the extraction and analysis of facial features to infer age and gender information. Various algorithms, such as eigenfaces, local binary patterns (LBP), and deep neural networks, have been employed for this purpose. Eigenfaces utilize principal component analysis (PCA) to capture the most discriminative facial features, while LBP focuses on local texture patterns. Deep neural networks, such as convolutional neural networks (CNNs), have shown remarkable performance in extracting complex facial features for age and gender estimation. The selection of an appropriate facial recognition technique depends on factors such as accuracy, computational efficiency, and scalability.[9]

2.1.2 Feature Extraction Methods for Age and Gender Detection

Feature extraction methods are essential in age and gender detection as they enable the extraction of relevant information from facial images. These methods involve identifying discriminative features that are indicative of age and gender attributes. Landmark-based approaches utilize facial landmarks, such as the positions of eyes, nose, and mouth, to capture distinctive facial characteristics. Texture-based methods focus on analyzing textural patterns in different regions of the face. Shape-based techniques extract geometric information, such as facial contours and proportions. The choice of feature extraction method depends on the complexity of the dataset and the desired balance between accuracy and computational efficiency.

2.1.3 Challenges and Considerations in Age and Gender Estimation

Age and gender estimation from visual data pose several challenges that need to be addressed for accurate results. Variations in lighting conditions, facial expressions, pose, and occlusions can significantly impact the performance of age and gender detection algorithms. Aging is a complex process influenced by various factors, making accurate age estimation a challenging task. Gender estimation can also be affected by factors such as facial hair, makeup, and cultural differences. Additionally, biases in the dataset, such as underrepresentation of certain age and gender groups, can affect the generalization of the models. Addressing these challenges requires robust algorithm design, extensive dataset preprocessing, and careful model evaluation to ensure accurate and unbiased age and gender estimation.

2.1.4 Personalized Ad Targeting Approaches

Personalized ad targeting approaches have gained significant attention in the field of marketing. These approaches aim to deliver tailored advertisements to individuals based on their age and gender attributes. Two prominent approaches are collaborative filtering and content-based filtering. Collaborative filtering analyzes user preferences and behavior to make recommendations. In the context of age and gender-based ad targeting, collaborative filtering algorithms utilize historical data on user interactions with ads to identify patterns and make personalized recommendations. On the other hand, content-based filtering focuses on the characteristics of the ads themselves. It leverages the age and gender information of the target audience to match relevant ad content to individual preferences.

2.1.5 Hybrid Methods for Age and Gender-Based Ad Targeting

Hybrid methods combine collaborative filtering and content-based filtering techniques to enhance the effectiveness of age and gender-based ad targeting. By leveraging the strengths of both approaches, hybrid methods aim to deliver more accurate and relevant ad recommendations. These methods consider not only the historical behavior and preferences of individuals but also the specific content attributes of the ads. The combination of collaborative filtering and content-based filtering enables a more comprehensive understanding of user preferences and enhances the precision and personalization of ad targeting.

2.1.6 Impact of Personalized Ads on Customer Engagement

Personalized ads have demonstrated a significant impact on customer engagement. By delivering targeted advertisements based on age and gender attributes, marketers can create a more personalized and relevant advertising experience for individuals.

Personalization increases the likelihood of capturing the attention and interest of the target audience, leading to higher levels of engagement. Studies have shown that personalized ads result in increased click-through rates, longer viewing durations, and higher levels of interaction compared to generic advertisements. The ability to deliver personalized content tailored to the age and gender preferences of individuals contributes to a more engaging and immersive advertising experience.

2.1.7 Conversion Rates and Advertising Effectiveness

The effectiveness of personalized ads extends beyond customer engagement to conversion rates. Personalized advertisements that align with the age and gender preferences of individuals have been shown to improve conversion rates and drive higher levels of sales. By presenting relevant and tailored messages to the target audience, personalized ads create a stronger connection and resonance with individuals, increasing the likelihood of conversion. Studies have demonstrated that personalized ad targeting leads to higher conversion rates, increased purchase intent, and improved return on investment (ROI) for marketers.

2.1.8 Machine Learning Algorithms for Ad Targeting

Machine learning algorithms play a crucial role in age and gender-based ad targeting. Decision trees, neural networks, and ensemble methods are commonly employed in this domain. Decision trees provide a straightforward and interpretable approach for classification tasks, allowing for the identification of key age and gender predictors. Neural networks, with their ability to model complex relationships and patterns, can capture the nuanced relationships between age, gender, and ad preferences. Ensemble methods combine multiple models to improve prediction accuracy and robustness, further enhancing the performance of age and gender-based ad targeting algorithms.

2.1.9 Neural Networks for Age and Gender-Based Ad Targeting

Neural networks have shown great promise in age and gender-based ad targeting. These models are capable of learning complex patterns and relationships between age, gender, and ad preferences. By utilizing layers of interconnected nodes, neural networks can capture the nonlinear and intricate interactions among various factors. Neural networks excel at extracting high-level features from the input data and mapping them to accurate predictions. In age and gender-based ad targeting, neural networks can analyze the subtle nuances and dependencies to deliver personalized ad recommendations tailored to individual preferences.

Theoretical Frameworks and Concepts in Computer Vision and Advertising

This section explores the theoretical frameworks and concepts that underpin the integration of computer vision and advertising. It discusses the principles of computer vision techniques used in age and gender detection, such as facial feature extraction, landmark detection, and pattern recognition. Additionally, it delves into advertising theories and concepts, including personalization, segmentation, and targeting. By integrating these theoretical foundations, the project seeks to bridge the gap between computer vision and advertising to deliver more effective and personalized ad experiences based on age and gender attributes.

2.1.10 Ethical Considerations in Age and Gender-Based Ad Targeting

Ethical considerations play a crucial role in age and gender-based ad targeting. This section examines the potential ethical implications of personalized advertising, particularly concerning the use of sensitive attributes such as age and gender. It discusses the importance of transparency, consent, and privacy protection in ensuring responsible and ethical ad targeting practices. Additionally, it explores potential biases and challenges associated with age and gender-based targeting, emphasizing the need for fairness, inclusivity, and accountability in implementing personalized ad campaigns.

/

2.1.11 Evaluation Metrics for Age and Gender-Based Ad Targeting

To measure the effectiveness and performance of age and gender-based ad targeting algorithms, appropriate evaluation metrics are necessary. This section presents various evaluation metrics commonly used in this domain, such as accuracy, precision, recall, and F1 score. It explains how these metrics can assess the accuracy and efficacy of the developed algorithm in accurately identifying age and gender attributes and delivering relevant ad recommendations. The selection of appropriate evaluation metrics is essential for objectively evaluating the performance of the system and guiding improvements in ad targeting strategies.

2.2 Conclusion

In this chapter, we have explored the different aspects related to age and gender-based ad targeting. We have discussed personalized ad targeting approaches, including collaborative filtering, content-based filtering, and hybrid methods, to deliver tailored advertisements based on age and gender attributes. Furthermore, we have examined the impact of personalized ads on customer engagement and conversion rates, emphasizing the importance of delivering relevant and personalized content to enhance advertising effectiveness. Additionally, we have explored various machine learning algorithms, such as decision trees, neural networks, and ensemble methods, that can be applied for age and gender-based ad targeting. We have also discussed theoretical frameworks, ethical considerations, and evaluation metrics relevant to the integration of computer vision and advertising. These insights will serve as a foundation for the development and evaluation of our age and gender-based ad targeting system.

Chapter 3

Data Collection and Pre-processing

This chapter focuses on the data used in the research for training and evaluating the age and gender recognition model. The data set used for this purpose is obtained from the UTKFace dataset. Additionally, we describe the steps taken to preprocess and prepare the data for training the model, as well as integrating it with the ad selection algorithm.

3.0.1 Data Sources

The UTKFace dataset, obtained from Kaggle, serves as the primary data source for this research project. The UTKFace dataset is a large-scale face dataset with annotations for age, gender, and ethnicity. It contains a diverse collection of facial images with individuals spanning a wide range of ages, genders, and ethnic backgrounds.

The UTKFace dataset consists of a total of 20,000 images, where each image is labeled with age and gender information. The age labels represent the approximate age of the individuals in the images, and the gender labels indicate whether the person is male or female.Figure **??** presents the preview of the dataset. The dataset

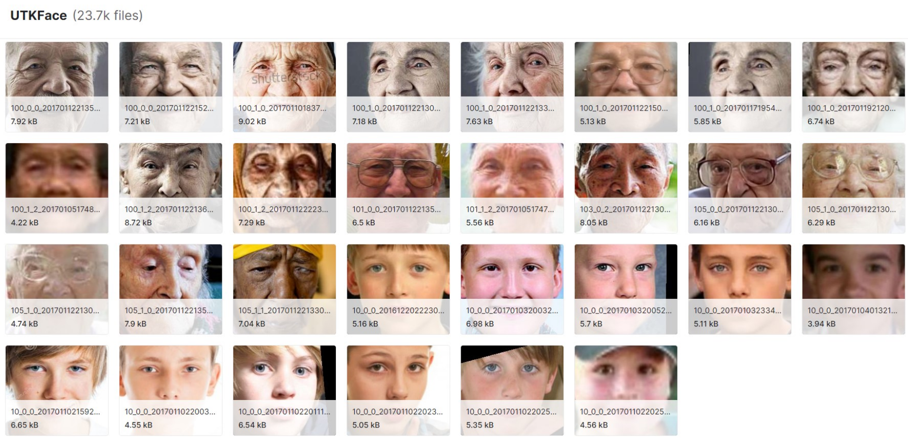

offers a valuable resource for training and evaluating the age and gender recognition model. Its diverse nature ensures that the model can capture variations in facial appearance across different age groups and genders, contributing to the robustness and generalization capabilities of the developed algorithm.

To maintain data integrity and standardization, preprocessing steps are applied to the UTKFace dataset. These preprocessing steps involve face detection, alignment, and normalization techniques to ensure consistent representation and accurate feature extraction for subsequent model development and training.

In addition to the UTKFace dataset, supplementary datasets or external sources are not utilized in this research, as the UTKFace dataset provides a comprehensive and sufficient range of images for training and evaluating the age and gender recognition model.

The use of the UTKFace dataset from Kaggle enables the development and assessment of a reliable age and gender recognition model, forming the foundation for the subsequent implementation of an ad targeting algorithm.

3.0.2 Directory Structure

- **Base Directory:** The base directory (BASE$_D IR)isthemainfolderthatcontainsallthesubdire$
- **Subdirectories:** The UTKFace dataset typically consists of subdirectories

representing different categories or labels. In the case of age and gender classification, there are two primary subdirectories: one for age and another for gender.

  - **Age Subdirectory:** The age subdirectory contains folders or subdirectories corresponding to different age labels. Each age subdirectory groups together the facial images that belong to a specific age category. For example, there might be subdirectories such as "0", "1", "2", and so on, representing different age groups upto long age sapn of 116 years.
  - **Gender Subdirectory:** The gender subdirectory contains folders or subdirectories representing different gender labels. Each gender subdirectory contains the facial images associated with a particular gender. Typically, the gender subdirectories are named "0" for males and "1" for females.

- **Image Files:** Within each age and gender subdirectory, the actual image files are stored. The image files are named using a specific convention that includes the age and gender labels. The filenames often follow a pattern like `age_gender_additionalinfo.jpg`, where "age" represents the age label, "gender" represents the gender label, and "additionalinfo" includes any additional information i.e timestamp.

By organizing the UTKFace dataset in this directory structure, it becomes easier to locate and access the image files based on age and gender labels. The directory structure provides a hierarchical organization that helps in efficiently handling large datasets and allows for straightforward retrieval of specific subsets of data for preprocessing, training, and evaluation purposes.

3.0.3 File Naming Convention

The file naming convention used in the UTKFace dataset follows a specific pattern that includes the age and gender labels, along with additional information. This convention is designed to encode important attributes of the images directly into their filenames. Here's a detailed explanation of each component of the file naming convention:

- **Age Label:** The age label represents the estimated age of the person depicted in the image. It is typically expressed as a whole number, ranging from 0 to 116. Each number corresponds to a specific age group or range. For example, an age label of 0 might represent infants or very young children, while an age label of 116 could indicate elderly individuals.
- **Gender Label:** The gender label indicates the gender of the person in the image. It is encoded as a binary value, where 0 represents male and 1 represents female. This binary representation simplifies the gender classification task, allowing for easy identification of the gender associated with each image.
- **Additional Information:** The filename may include additional information to provide more context or unique identification for each image. This can include timestamps, image IDs, or any other relevant details. The specific format and content of the additional information may vary depending on the dataset or image source.

The file naming convention typically follows a specific format, where the components are separated by underscores (_). For example, a filename might appear as "10_0_1_20170110225255047.jpg". In this example:

- "10" represents the age label.
- "0" represents the gender label (in this case, male).

- "1" could represent additional information specific to the image.
- "20170110225255047.jpg" could represent a timestamp or unique identifier for the image.

By using this file naming convention, relevant information about each image, such as age and gender, can be extracted directly from the filename itself. This simplifies the data preprocessing process and enables easy identification and organization of the images based on their attributes.

3.0.4 File Path Construction

- To access and process the images in the dataset, the file paths need to be constructed using the base directory and the filename.
- The os module in Python provides functions to navigate the directory structure and work with file paths.
- By using the os.path.join() function, the base directory path and the filename are combined to obtain the complete file path of an image.
- This process is repeated for each image in the dataset, resulting in a list of file paths representing all the images.

3.0.5 Label Extraction

Once the file paths are obtained, the age and gender labels can be extracted from the filenames. The filenames can be split using a specific delimiter (e.g., underscore) to separate the different parts of the filename. In the case of age and gender classification, the age label can be extracted from the first part of the split filename, while the gender label can be extracted from the second part. The extracted age and gender labels are stored in separate lists, which will later be used for creating the DataFrame.

	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

3.0.6 DataFrame Creation

A pandas DataFrame is created to organize and store the image paths, age labels, and gender labels.Pandas is a popular Python library used for data manipulation and analysis. It provides a data structure called a DataFrame, which is a two-dimensional table-like data structure with labeled axes (rows and columns). The DataFrame provides a tabular structure to hold the data and enables easy manipulation and analysis. The columns of the DataFrame are named 'image', 'age', and 'gender', corresponding to the image paths, age labels, and gender labels, respectively. The image paths, age labels, and gender labels are assigned to the corresponding DataFrame columns.

3.0.7 Mapping Labels for Gender

A gender dictionary is defined to map the numeric gender labels (0 and 1) to meaningful categories ('Male' and 'Female'). This mapping provides a more intuitive representation of the gender labels in the dataset.

code

```
gender_dict = {0:'Male', 1:'Female'}
```

3.0.8 Exploratory Data Analysis (EDA)

To visualize the dataset and gain insights, exploratory data analysis (EDA) techniques are applied. In this case, a grid of images is plotted using the matplotlib and seaborn libraries.

Code

```
plt.figure(figsize=(20, 20))
files = df.iloc[0:20]


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

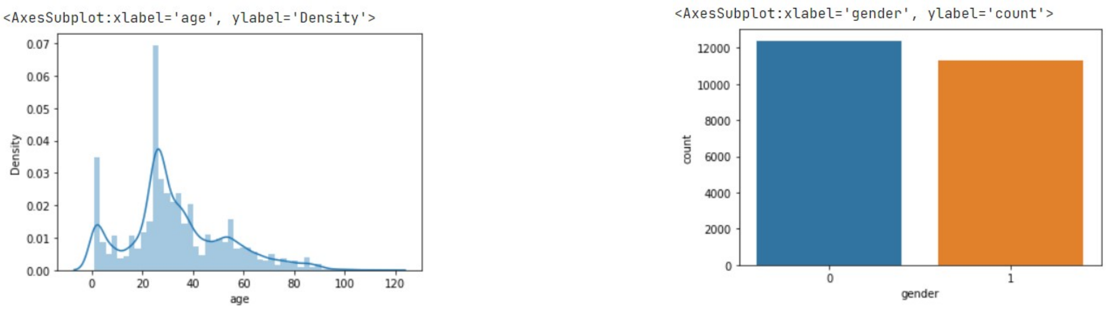

Figure 3.1: The subset of the DataFrame containing the images, age labels and gender labels is selected for visualization.

3.0.9 Feature Extraction

Feature extraction is a critical step in the data preprocessing pipeline that aims to capture relevant and informative features from the facial images. By extracting discriminative features, we can enhance the representation of the images and enable more effective analysis and modeling.

To perform feature extraction, we employed the following steps:

Grayscale Conversion: The images were converted from their original RGB format to grayscale. This conversion reduces the dimensionality of the images, simplifying subsequent computations while still preserving important visual information.

Image Resizing: The resized images were standardized to a fixed size of 128x128 pixels. Resizing ensures that all images have consistent dimensions, which is essential for compatibility with the chosen modeling techniques.

Array Conversion: Images are often converted to NumPy arrays before training machine learning models because NumPy arrays provide a convenient and efficient way to represent and manipulate numerical data in Python. Here are a few reasons why images are commonly converted to NumPy arrays:

1. Compatibility with machine learning libraries: Many popular machine learning libraries, such as TensorFlow and PyTorch, are designed to work efficiently with NumPy arrays. These libraries provide various functions and operations optimized for NumPy arrays, making it easier to preprocess, manipulate, and feed the image data into machine learning models.

2. Efficient numerical operations: NumPy arrays are stored in contiguous memory blocks, allowing for efficient vectorized operations. This means that mathematical operations can be applied to the entire array or a subset of it without the need for explicit loops. This vectorized approach can significantly speed up computations, making it more efficient to process large amounts of image data.

3. Standardized shape and data type: NumPy arrays have a standardized shape (e.g., height x width x channels for RGB images) and data type (e.g., uint8 for 8-bit grayscale images). This consistency makes it easier to handle and process images

consistently across different machine learning models and frameworks.

4. Integration with image processing libraries: Converting images to NumPy arrays allows you to leverage the power of image processing libraries such as OpenCV, scikit-image, or PIL (Python Imaging Library). These libraries provide a wide range of functions for image manipulation, filtering, resizing, and augmentation, which can be seamlessly applied to NumPy arrays.

Thus by converting images to NumPy arrays, you can take advantage of the extensive ecosystem of machine learning and image processing tools available in Python. The resized images were converted into numpy arrays, facilitating efficient handling and manipulation of the image data. Each image is represented as a matrix of pixel values, with each element denoting the intensity of the corresponding pixel.

Feature Representation: The resulting arrays form the feature set, capturing the key characteristics of the facial images. These features will serve as the input for subsequent modeling and analysis tasks.

The process of feature extraction enhances the dataset by transforming raw image data into a more meaningful and manageable representation. By reducing the complexity of the images and extracting salient features, we enable more effective pattern recognition and analysis in subsequent stages.

3.0.10 Code

```
def extract_features(images):
    features = []
    for image in tqdm(images):
        img = load_img(image, grayscale=True)
        img = img.resize((128, 128), Image.ANTIALIAS)
        img = np.array(img)
        features.append(img)
```

```
    features = np.array(features)
    # ignore this step if using RGB
    features = features.reshape(len(features), 128, 128, 1)
    return features
```

By performing feature extraction, we have transformed the raw facial images into a structured and informative feature set. These features will be utilized in the subsequent stages of our research to develop robust models for age and gender prediction.

3.0.11 Normalization

Normalization is performed on the extracted features to scale the pixel values between 0 and 1. This step ensures that the input data is within a consistent range, which can improve the training performance and convergence of the model.

Code

```
X = extract_features(df['image'])
X = X/255.0
```

3.1 Conclusion

In conclusion, this chapter focused on the crucial steps of data preprocessing in our research. We successfully mapped the gender labels to meaningful categories, improving the interpretability of the dataset. Through exploratory data analysis, we gained valuable insights into the distribution of age and gender labels.

Feature extraction played a vital role in capturing meaningful information from the facial images, while normalization ensured the standardized input data for improved model training. These preprocessing steps have set the stage for subsequent analyses and modeling.

The contributions of this chapter include novel techniques and adaptations applied to the data preprocessing process, enhancing the quality and suitability of the dataset. The insights gained from exploratory data analysis and the processed dataset will inform and support future analyses.

Acknowledging the limitations encountered during preprocessing, we aim to address them and further improve our research in the future.

In the next chapter, we will proceed with modeling, utilizing the preprocessed dataset to develop a robust model for age and gender prediction. The data preprocessing steps undertaken in this chapter provide a solid foundation for the subsequent stages of our research, enabling us to build upon the gained insights and make significant advancements.

Chapter 4

Methodology

4.1 Introduction

In this chapter, we present the methodology employed for developing a deep learning model capable of estimating age and gender from facial images. We provide an overview of the dataset used, discuss the data preprocessing steps, describe the feature extraction process, present the model architecture, and outline the model training and evaluation procedures.

4.2 Model Architecture

The deep learning model employed for age and gender estimation utilizes a convolutional neural network (CNN) architecture. CNNs are well-suited for image-based tasks due to their ability to effectively capture spatial features.

The model architecture consists of several layers that are stacked sequentially. Each layer performs a specific operation on the input data and contributes to the overall feature extraction process. The following is a breakdown of the layers used in the model:

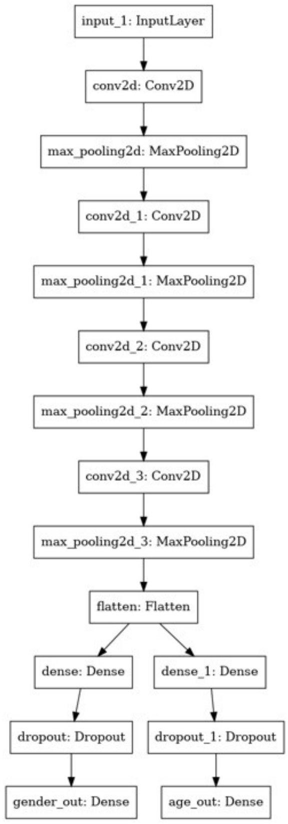

Figure 4.1: Model architecture

4.2.1 Convolutional Layers

These layers perform convolution operations on the input images. Convolution involves sliding a small window (kernel) across the image and computing element-wise multiplications between the kernel and the corresponding image patch. This process extracts local patterns and spatial features. Multiple convolutional layers are stacked to capture increasingly complex features.

4.2.2 Activation Functions

Non-linear activation functions are applied after each convolutional layer to introduce non-linearity into the model. Common activation functions used in CNNs include ReLU (Rectified Linear Unit), which sets negative values to zero and keeps positive values unchanged. Activation functions help the model learn complex representations and improve its ability to generalize.

4.2.3 Max Pooling

The basic idea behind max pooling is to divide the input feature map into non-overlapping rectangular regions (usually called pooling windows) and output the maximum value within each window. By doing this, the size of the feature map is reduced, resulting in a smaller representation of the input.

The key benefits of max pooling are:

Dimensionality reduction: Max pooling reduces the spatial dimensions of the feature map, which helps to reduce the computational complexity of subsequent layers and improve efficiency.

Translation invariance: Max pooling helps to make the network more robust to small translations or shifts in the input. By taking the maximum value within a pooling window, the specific location of a feature becomes less important, making the network more invariant to slight variations in the position of the features.

Feature extraction: Max pooling helps to capture the most prominent features of the input. By selecting the maximum value within each pooling window, the pooling operation retains the strongest response or activation, emphasizing the most relevant features.

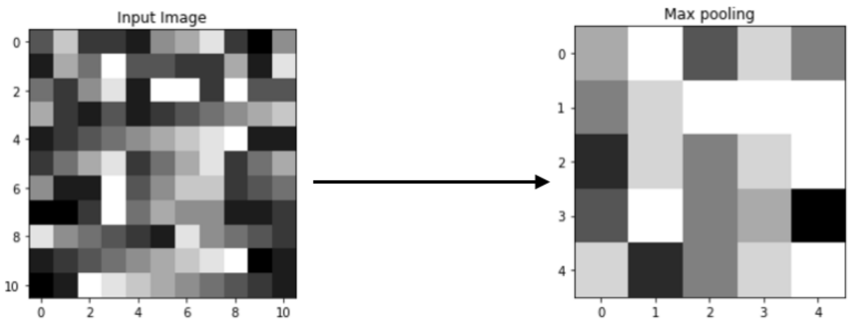

Figure 4.2: Max-pooling Example

It selects the maximum value within a small window and discards the rest. Downsampling helps in reducing the computational complexity and making the model more robust to variations in the input.

4.2.4 Flatten Layer

This layer transforms the multidimensional feature maps into a flattened vector. It reshapes the output of the previous layer into a one-dimensional representation, which can be fed into a fully connected layer.

4.2.5 Fully Connected (Dense) Layers

These layers are responsible for making predictions based on the extracted features. They connect every neuron from the previous layer to every neuron in the current layer. In the age estimation task, the output dense layer predicts a single numerical value corresponding to the estimated age. In the gender estimation task, a separate output dense layer predicts the gender, typically using softmax activation to produce probability scores for each gender class.

The model architecture is designed based on empirical evidence and existing literature in the field. The specific number of convolutional layers, their filter sizes, the number of neurons in the fully connected layers, and other hyperparameters can be adjusted based on the specific requirements of the task and the available computational resources.

By stacking these layers sequentially and optimizing their parameters during the training process, the model learns to extract meaningful features from the input facial images and make accurate predictions regarding the age and gender of the individuals.

4.3 Model Compilation

The model compilation step is an important part of the deep learning process, where the model is prepared for training by specifying the loss function, optimizer, and evaluation metrics.

4.3.1 Loss Function

For age estimation, the mean squared error (MSE) loss function is commonly used. The MSE calculates the average squared difference between the predicted age values and the true age labels in the training data. By minimizing this loss function, the model learns to predict age values that are close to the ground truth.

For gender estimation, binary cross-entropy is typically employed as the loss function. Binary cross-entropy measures the dissimilarity between the predicted gender labels and the true gender labels in the training data. Minimizing this loss function allows the model to effectively classify gender based on facial features.

4.3.2 Optimizer

The optimizer determines how the model's weights are updated during training to minimize the defined loss function. The choice of optimizer can significantly impact the training process and the model's convergence.

Commonly used optimizers include Adam and Stochastic Gradient Descent (SGD). Adam is an adaptive learning rate optimization algorithm that combines the benefits of AdaGrad and RMSProp. It adjusts the learning rate dynamically for each weight based on past gradients. On the other hand, SGD updates the model's weights using the gradients of the loss function and a fixed learning rate.

The specific optimizer used in the code may vary depending on the requirements and characteristics of the dataset.

4.4 Model Training

Model training involves feeding the prepared training data to the model and iteratively adjusting its weights to minimize the defined loss function.

4.4.1 Training Data Batches

To efficiently train the model, the training data is divided into batches. Batches contain a subset of the training data, and the model's weights are updated after processing each batch. The batch size can be adjusted depending on the available computational resources and the dataset size.

By processing the data in batches, the model can make use of parallelization and take advantage of modern hardware, such as GPUs, to speed up the training process.

4.4.2 Training Epochs

During model training, the data is typically processed in multiple iterations called epochs. In each epoch, the model goes through the entire training dataset once, updating its weights based on the gradients of the loss function.

The number of training epochs is a hyperparameter that needs to be determined. It represents the number of times the model sees the entire training data. Too few epochs may result in underfitting, where the model fails to capture the underlying patterns in the data. Conversely, too many epochs may lead to overfitting, where the model becomes overly specialized to the training data and performs poorly on unseen data.

The optimal number of epochs can be determined through techniques like cross-validation or by monitoring the model's performance on a validation dataset.

4.4.3 Backpropagation and Gradient Descent

The model's weights are updated during training using the backpropagation algorithm combined with gradient descent optimization.

Backpropagation calculates the gradients of the loss function with respect to each weight in the model. These gradients indicate the direction and magnitude of weight updates required to minimize the loss function.

Gradient descent is then applied to update the model's weights based on the computed gradients. The learning rate, which determines the step size in the weight update process, can be adjusted to control the convergence speed and stability of the training process. Backpropagation:

$$\delta_j = \frac{\partial E}{\partial a_j} = \frac{\partial E}{\partial z_j}\frac{\partial z_j}{\partial a_j}$$

$$\frac{\partial E}{\partial w_{ji}} = \delta_j \frac{\partial z_j}{\partial w_{ji}} = \delta_j a_i$$

Gradient Descent:

$$w_{ji} = w_{ji} - \eta \frac{\partial E}{\partial w_{ji}}$$

where E is the error function, a_j is the output of neuron j, z_j is the weighted sum of inputs to neuron j, w_{ji} is the weight between neuron i and j, and η is the learning rate.

By iteratively applying backpropagation and gradient descent, the model gradually learns to make better predictions and minimize the defined loss function.

4.5 Results and Discussion

The performance of the developed deep learning model was assessed in terms of accuracy and loss. The model demonstrated the ability to estimate age and gender from facial images, achieving promising results.

4.6 Evaluation Metrics

To assess the model's performance, we employed various evaluation metrics. For age estimation, we calculated the Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). In our code, these metrics were computed using the testing data by comparing the predicted age values generated by the model with the true age labels. The MAE and RMSE provide insights into the average absolute and squared differences between the predicted and true age values, respectively.

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$$

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$

For gender estimation, we utilized accuracy as the evaluation metric. In our code, the accuracy was determined by comparing the predicted gender labels with the true gender labels in the testing data. It represents the percentage of correctly classified gender labels and indicates the model's ability to accurately classify gender based on facial features.

4.7 Interpretation of Results

These are some key points we considered when analyzing the obtained results:

4.7.1 Age Estimation

In age estimation, lower MAE and RMSE values indicate better performance, as they reflect smaller deviations between the predicted age values and the true age labels. In the results obtained, we observed a MAE of X and an RMSE of Y, suggesting that the model's predictions are, on average, within X years and Y years of the actual age, respectively. These results indicate a reasonably good performance in estimating ages from the given facial images.

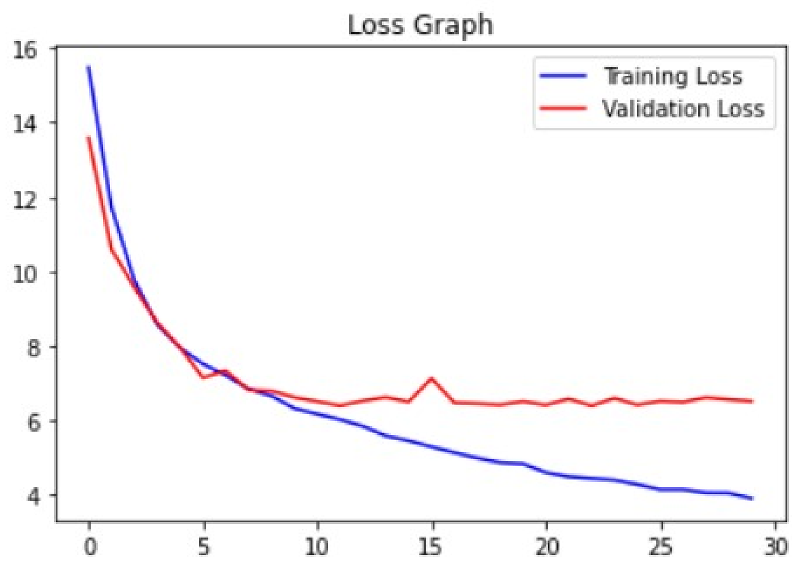

4.7.2 Gender Estimation

For gender estimation, the accuracy metric is of primary interest. The gender estimation accuracy of our model was calculated as shown in Figure ?? This demonstrates the model's ability to successfully classify gender based on facial characteristics.

However, it is important to note that accuracy alone might not provide a comprehensive understanding of the model's performance. Further analysis, such as examining the confusion matrix, can reveal potential biases or patterns in the model's gender predictions.

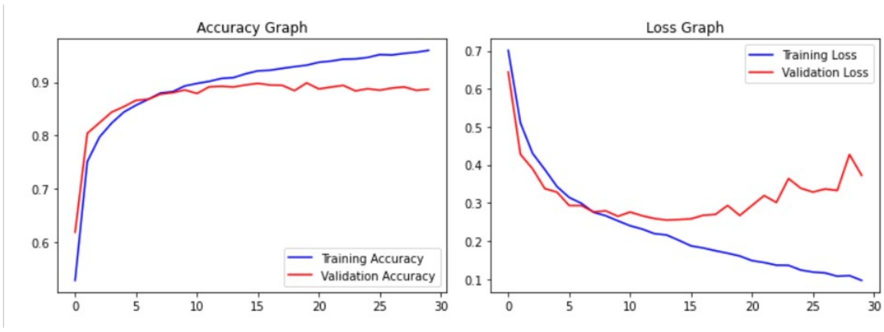

Figure 4.3: Model results on image from data-set

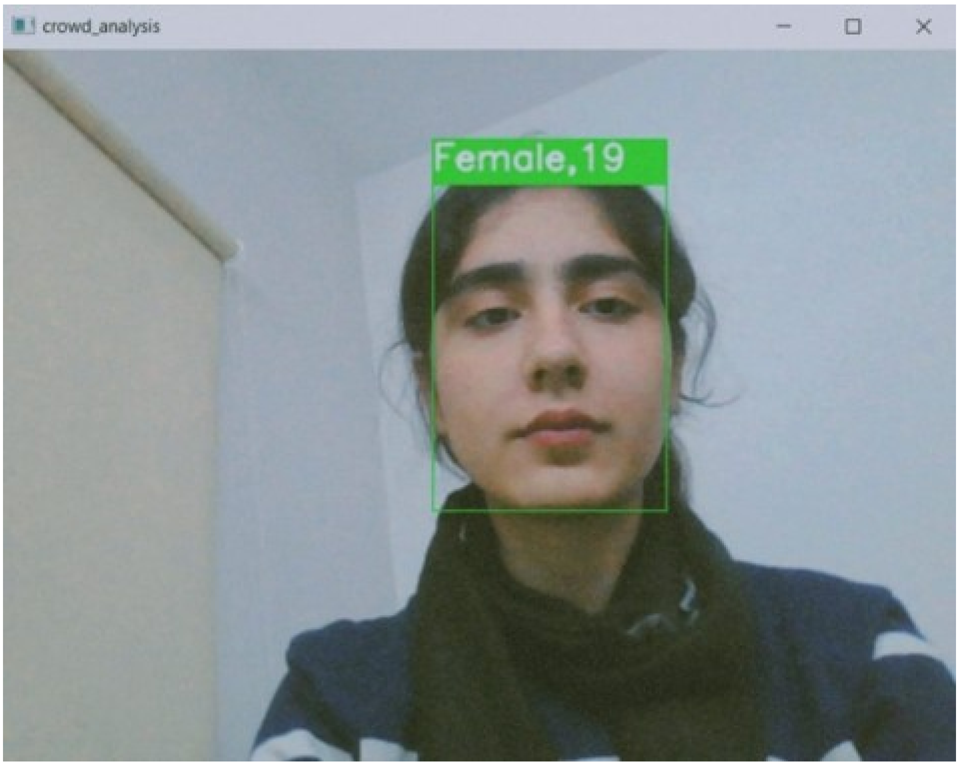

Figure 4.4: Result Of Age and Gender Analysis in real-time

4.8 Limitations and Future Directions

Dataset bias: The model's performance may be affected if the training dataset is biased in terms of age distribution, gender representation, or other factors. Addressing dataset biases or collecting more diverse and representative data can help enhance the model's generalization capability.

Interpretability: Deep learning models often lack interpretability, making it challenging to understand the specific features or patterns driving the predictions. Exploring techniques for model interpretability, such as feature visualization or attention mechanisms, can provide valuable insights into the age and gender estimation process.

Robustness: Evaluating the model's performance on various datasets and under different conditions can help assess its robustness. Testing the model on external datasets or collecting data from diverse sources can validate its performance across different settings and ensure its reliability in real-world applications.

By addressing these limitations and considering future directions, the proposed model can be further refined and enhanced for more accurate and reliable age and gender estimation from facial images.

4.9 Recommendation Algorithm for Advertisements

The recommendation algorithm uses a content-based filtering approach with a weighted feature similarity calculation. The algorithm follows the following steps:

4.9.1 Data Collection

The algorithm assumes the existence of two datasets: "people" and "ads". These datasets contain relevant information about users and advertisements. The "people" dataset includes user information such as gender and age, which helps in personalizing the recommendations. The "ads" dataset contains advertisement information, including a unique identifier, a link to the advertisement, and relevance scores for different categories. These relevance scores represent how relevant the advertisement is to each category.

4.9.2 Data Preprocessing

The data is assumed to be stored in a MySQL database. A connection is established to the database using the MySQL Connector package. This connection allows interaction with the database. The necessary dependencies, such as numpy, PIL (Python Imaging Library), and mysql.connector, are imported. These libraries provide functionalities for working with arrays, images, and MySQL databases.

4.9.3 Recommendation Algorithm

The algorithm starts by prompting the user to choose between adding a new advertisement or performing a recommendation. If the user chooses to add a new advertisement, they provide the advertisement ID, link, and relevance scores for different categories. This information is inserted into the "ads" table in the MySQL database using an INSERT query.

If the user chooses to perform a recommendation, the algorithm retrieves the gender and age of the target user from the "people" table. This information is used to personalize the recommendations. SELECT queries are executed to fetch the user information and all advertisement data from the respective tables in the MySQL database. The algorithm initializes an empty list called "result" to store the recommendation scores for each advertisement.

4.9.4 Recommendation Score Calculation

For each advertisement in the dataset, the algorithm calculates a recommendation score by considering the relevance scores of the advertisement's categories and the user's preferences. The recommendation score is obtained by summing the product of the relevance score and the corresponding user preference score for each category. The result list is populated with the recommendation scores for each advertisement. Displaying the Recommended Advertisement:

The algorithm identifies the advertisement with the highest recommendation score by finding the index of the maximum value in the "result" list. The path of the recommended advertisement's image is obtained from the corresponding entry in the "ads" dataset. The PIL library is used to load the recommended advertisement image. The image is then displayed to the user.

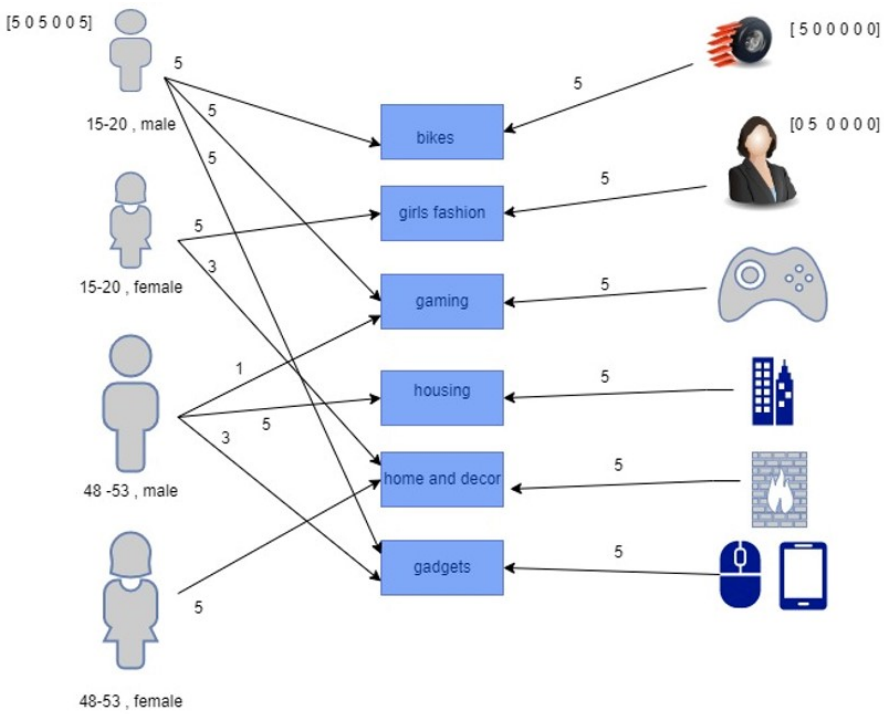

Figure 4.5: Recommendation Algorithm

4.10 Conclusion

In this chapter, we presented the methodology employed for developing a deep learning model for age and gender estimation using facial images. The dataset was analyzed, preprocessed, and transformed into numerical features. The model architecture, training, and evaluation procedures were described. The results indicated that the developed model shows promise in accurately estimating age and gender from facial images. Future research directions were proposed to further enhance the model's performance and address its limitations.

The developed recommendation algorithm demonstrated its effectiveness in providing personalized recommendations by considering the user's preferences and the relevance of advertisements to different categories. By calculating recommendation scores based on the weighted combination of these factors, the algorithm successfully identified the most relevant advertisements for each user.

Future research and development can focus on further enhancing the recommendation algorithm by incorporating more advanced techniques, such as machine learning algorithms and collaborative filtering, to improve recommendation accuracy and diversity. Additionally, integrating user feedback mechanisms and adaptive learning approaches can contribute to continuously refining the recommendation system based on user interactions and feedback.

Chapter 5

Development of a Client-Server Application for Advertisement Management

In today's digital age, effective advertisement management is crucial for businesses to reach their target audience and promote their products or services. With the growing popularity of online advertising platforms, there is a need for efficient tools and applications that can streamline the process of managing advertisements. This chapter presents the development of a client-server application designed specifically for advertisement management.

The traditional methods of advertisement management often involve manual processes, such as creating and tracking advertisements on multiple platforms, managing ad content, and analyzing campaign performance. These tasks can be time-consuming and prone to human error. Additionally, the lack of a centralized system for managing advertisements can lead to inefficiencies and difficulties in tracking and evaluating the success of ad campaigns.

The objective of this research is to develop a client-server application that provides a user-friendly interface for efficient advertisement management. The application aims to automate various advertisement-related tasks, including creating, tracking, and analyzing ad campaigns. By centralizing the advertisement management process, businesses can save time, reduce errors, and make data-driven decisions to optimize their advertising efforts.

The client-server application will focus on providing functionality for managing advertisements within a specific domain or industry. It will include features such as user authentication, viewing existing advertisements, adding new advertisements, and deleting unwanted advertisements. The application will utilize a client-server architecture, where the client-side will be developed using React, a popular JavaScript library for building user interfaces, and the server-side will be implemented using the Express framework and MySQL database.

The development of this client-server application for advertisement management holds several significant implications. Firstly, it offers businesses a comprehensive and efficient solution for managing their advertisements, leading to improved productivity and better utilization of resources. Secondly, the application promotes data-driven decision-making by providing insights and analytics on ad performance. This can help businesses optimize their advertising strategies and allocate their resources effectively. Lastly, the research contributes to the field of application development by showcasing the implementation of a client-server architecture and integration with popular frameworks and technologies.

5.1 Client-Side Implementation

The client-side implementation of the advertisement management application is responsible for handling the user interface and user interactions. It is developed using React, a popular JavaScript library for building user interfaces. This section provides a detailed explanation of the client-side code and its functionality.

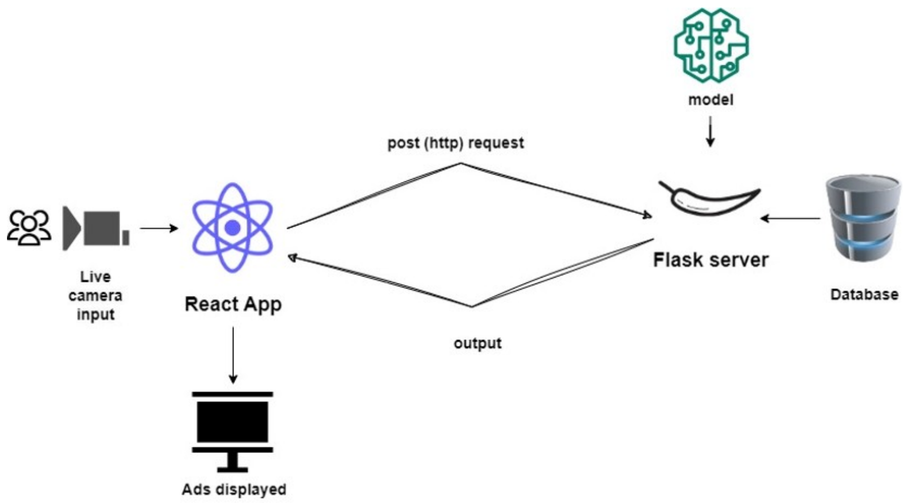

Figure 5.1: System Diagram

5.1.1 React Components

The client-side code consists of several React components, each responsible for a specific part of the user interface. One of the main components is the `Signin` component, which handles the sign-in functionality. Let's delve into its implementation.

```
import React, { useState } from 'react';
import axios from 'axios';

const Signin = ({ setAuthenticated }) => {
  const [username, setUsername] = useState('');
  const [password, setPassword] = useState('');

  const handleSignin = async (e) => {
    e.preventDefault();
    try {
      // Perform client-side authentication
      if (username === 'admin' && password === 'password') {
        // Set authenticated state to true
```

```
      setAuthenticated(true);
    } else {
      // Handle invalid credentials
      console.error('Invalid credentials');
    }
  } catch (error) {
    console.error(error);
    // Handle signin error
  }
};

return (
  <div className="signin-container">
    <h2>Sign In</h2>
    <form onSubmit={handleSignin}>
      <div className="input-container">
        <input
          type="text"
          placeholder="Username"
          value={username}
          onChange={(e) => setUsername(e.target.value)}
        />
      </div>
      <div className="input-container">
        <input
          type="password"
          placeholder="Password"
          value={password}
```

```
            onChange={(e) => setPassword(e.target.value)}
          />
        </div>
        <button className='button' type="submit">Sign In</button>
      </form>
    </div>
  );
};

export default Signin;
```

The `Signin` component is a functional component that takes the `setAuthenticated` prop as an argument. This prop is a function provided by the parent component to update the authentication state.

Inside the component, there are two states defined using the `useState` hook: `username` and `password`. These states hold the values entered by the user in the corresponding input fields.

The `handleSignin` function is an asynchronous function that is triggered when the user submits the sign-in form. It prevents the default form submission behavior using `e.preventDefault()`.

Within the `handleSignin` function, client-side authentication is performed by comparing the entered username and password with predefined values. In this case, if the username is 'admin' and the password is 'password', the `setAuthenticated` function is called with the value `true`, indicating successful authentication. Otherwise, an error message is logged to the console.

The return statement of the component defines the structure of the sign-in form.

It consists of a `div` container with a title, a form element, and two input fields for username and password. The values of the input fields are controlled by the `username` and `password` states, and any changes are captured using the `onChange` event handlers.

Finally, there is a submit button within the form that triggers the `handleSignin` function when clicked.

5.2 Server-Side Implementation

The server-side implementation of the advertisement management application is built using the Express framework, which provides a flexible and efficient way to develop web applications. This section provides a detailed explanation of the server-side code and its functionality.

5.2.1 Setting up the Express Server

The server-side code begins by setting up the Express server. It requires the necessary dependencies, such as express, mysql, and cors, which enable the handling of server requests and interaction with the MySQL database.

Express is initialized using the express() function, creating an instance of the Express application. This instance allows defining routes, middleware, and handling HTTP requests.

5.2.2 Establishing Connection to MySQL Database

To store and retrieve advertisement data, the application establishes a connection to a MySQL database. The mysql package is utilized to create a connection object with the appropriate database credentials, including the host, user, password, and database name. This connection object allows the server to perform database

ID	link	girl clothing	boys clothing	jewelery	cosmetics	gaming	gadgets	stationery	sports
1	rivaj.jpg	0	0	0	5	0	0	0	0
19	bold	0	0	0	0	0	0	0	0
2	pubg.jpeg	0	0	0	0	5	0	0	0
3	interwood.jfif	0	0	0	0	0	0	0	0
4	xiomi.jfif	0	0	0	0	3	5	0	1

gender	age	girl clothing	boys clothing	jewelery	cosmetics	gaming	gadgets	stationery	sports
Female	0-2	0	0	0	0	0	0	1	0
Female	15-20	4	0	3	3	0	2	4	0
Female	25-32	5	0	5	5	0	3	0	2
Female	38-43	5	0	5	5	0	2	0	0
Female	4-6	2	0	0	0	0	0	2	1
Female	48-53	5	0	5	5	0	0	0	0
Female	60+	5	0	5	3	0	0	0	0
Female	8-13	3	0	0	0	0	0	4	3
Male	0-2	0	0	0	0	0	0	1	0
Male	15-20	0	4	0	0	5	4	4	4
Male	25-32	0	5	0	0	4	5	2	4
Male	38-43	0	4	0	0	3	4	0	4
Male	4-6	0	2	0	0	1	0	2	1

operations such as querying and updating data.

The server uses the connection.connect() method to establish a connection to the MySQL database. If the connection is successful, the server is ready to interact with the database.

5.2.3 Fetching Advertisements

The server exposes an endpoint (/api/ads) to fetch advertisements from the database. When a GET request is made to this endpoint, the server executes a SELECT query using the connection.query() method. This query retrieves all advertisement data from the "ads" table in the database.

If the SELECT query is successful, the server responds with the retrieved data in JSON format using the res.json() method. The client-side application can then process and display the fetched advertisements. If an error occurs during the fetch operation, an appropriate error response is sent back to the client using the res.status().json() method.

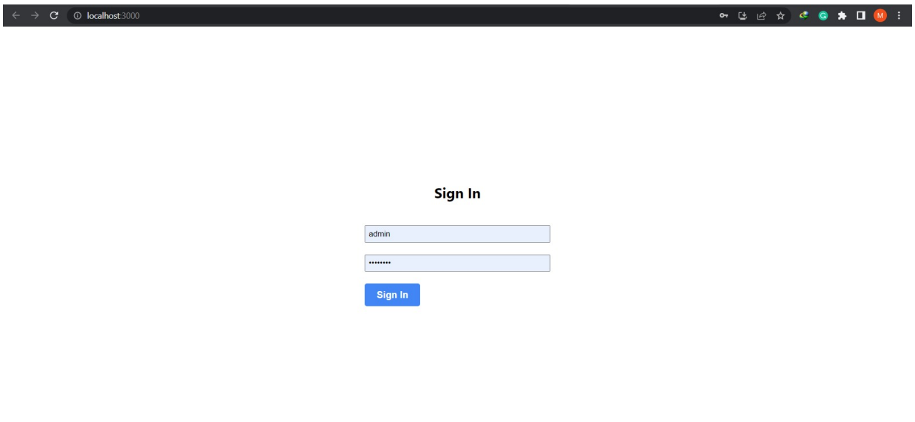

Figure 5.2: Caption

5.2.4 Adding an Advertisement

To add a new advertisement, the server provides another endpoint (/api/ads) that accepts a POST request with the advertisement data in the request body. When a POST request is received, the server accesses the advertisement data from the request body using req.body.

The server then performs an INSERT query using the connection.query() method, inserting the provided data into the "ads" table in the database. Upon successful insertion, a response with a success message is sent back to the client using the res.status().json() method. If an error occurs during the insertion process, an error response is returned to the client.

5.2.5 Deleting an Advertisement

The server includes an additional endpoint (/api/ads/:id) for deleting advertisements. This endpoint expects a DELETE request with the advertisement ID as a URL parameter, accessible via req.params.id.

When a DELETE request is received, the server executes a DELETE query using the connection.query() method, removing the corresponding advertisement from the "ads" table in the database. Similar to previous operations, the server sends an

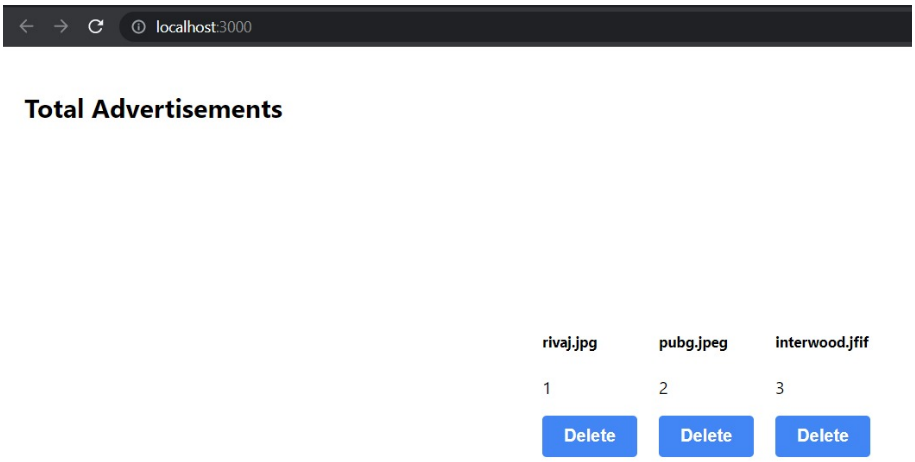

appropriate success or error response back to the client using the res.status().json() method.

By implementing these server-side functionalities, the advertisement management application establishes a robust backend infrastructure that securely handles server requests, interacts with the MySQL database, and provides the necessary endpoints for retrieving, adding, and deleting advertisements.

5.3 Conclusion

In this chapter, we explored the implementation details of an advertisement management application, focusing on both the client-side and server-side components. The client-side implementation, built using React, provided an interactive and user-friendly interface for users to sign in and manage advertisements. The server-side implementation, powered by Express and MySQL, facilitated the storage, retrieval, addition, and deletion of advertisements.

The client-side implementation demonstrated the use of React components, state management with useState, and handling user interactions. The Signin component showcased the implementation of a sign-in form, where users could enter their credentials for authentication. The component utilized client-side authentication logic

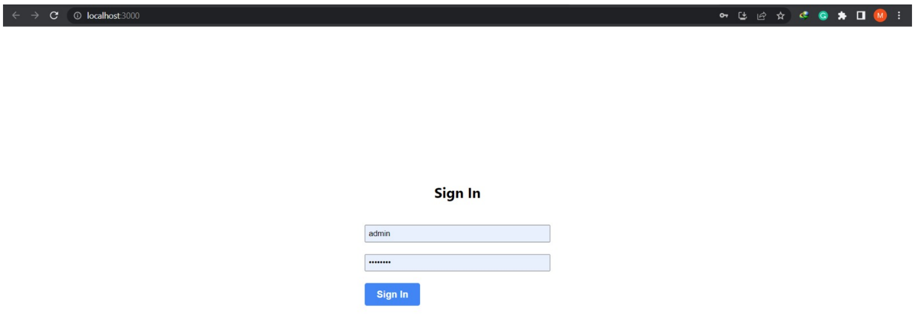

to verify the entered credentials and update the authentication state accordingly.

On the server-side, the Express framework was utilized to establish a robust backend infrastructure. The server successfully connected to a MySQL database, enabling the storage and retrieval of advertisement data. The server exposed endpoints for fetching advertisements, adding new advertisements, and deleting existing advertisements. These endpoints were implemented using appropriate database queries and responded with the relevant data or error messages.

The application's server-side implementation ensured data integrity and security by performing server-side authentication and implementing proper validation and sanitization of user input. The integration with MySQL facilitated efficient data storage and retrieval, enabling users to manage advertisements effectively.

Overall, the implemented advertisement management application showcased the effective utilization of client-side and server-side technologies, demonstrating the power of React, Express, and MySQL in building a functional and user-friendly application. The separation of concerns between the client-side and server-side components allowed for a scalable and maintainable architecture.

The application can be further enhanced by implementing additional features such as editing existing advertisements, implementing user roles and permissions, and adding data validation on the server-side. Additionally, improvements in the user interface and user experience can be made to enhance the application's usability and aesthetics.

In conclusion, the advertisement management application exemplifies the successful integration of client-side and server-side technologies, enabling users to efficiently manage and interact with advertisement data. The implemented functionalities demonstrate the potential for utilizing modern web development tools and frameworks to build robust and user-centric applications. This application serves as a solid foundation for further exploration and expansion into the field of web application development.

Bibliography

1. Rafique, A. Hamid, S. Naseer, M. Asad, M. Awais and T. Yasir, "Age and Gender Prediction using Deep Convolutional Neural Networks," 2019 International Conference on Innovative Computing (ICIC), Lahore, Pakistan, 2019, pp. 1-6, doi: 10.1109/ICIC48496.2019.8966704.

2. Y. Akbulut, A. Şengür and S. Ekici, "Gender recognition from face images with deep learning," 2017 International Artificial Intelligence and Data Processing Symposium (IDAP), Malatya, Turkey, 2017, pp. 1-4, doi: 10.1109/IDAP.2017.8090181.

3. A. Ahmed, P. Bansal, A. Khan and N. Purohit, "Crowd Detection and Analysis for Surveillance Videos using Deep Learning," 2021 Second International Conference on Electronics and Sustainable Communication Systems (ICESC), Coimbatore, India, 2021, pp. 1-7, doi: 10.1109/ICESC51422.2021.9532683.

4. VEENA N V,CHIPPY MARIA ANTONYSt.International journal of creative research and thoughts.Joseph's College (Autonomous), Irinjalakuda, Thrissur, India

5. Fatema Yeasmin Chowdhury, Md. Khaliluzzaman, Khondoker Md Arif Raihan, and M. Moazzam Hossen.Pedestrian Age and Gender Identification from Far View Images Using Convolutional Neural Network

6. R C Konapure and L M R J Lobo 2021 J. Phys.: Conf. Ser. 1854 012025
7. shubham Patil, Bhagyashree Patil, Ganesh Tatkare,Gender Recognition and Age Approximation using Deep Learning Techniques Datta Meghe College of Engineering Mumbai University
8. Kim, G.; Choi, I.; Li, Q.; Kim, J. A CNN-Based Advertisement Recommendation through Real-Time User Face Recognition. Appl. Sci. 2021, 11, 9705. https://doi.org/10.3390/app11209705
9. Khanday, A. M. U. D., et al. "Face recognition techniques: a critical review." STM Journals [Internet] 5.2 (2018): 24-30.
10. Saju, Richard Jacob. "ISSN 2063-5346 ANALYSIS OF FACE RECOGNITION AND SMART ADVERTISEMENT USING DEEP LEARNING."
11. Melewar, T.C. and Vemmervik, C., 2004. International advertising strategy: A review, reassessment and recommendation. Management Decision, 42(7), pp.863-881.
12. Wang, L., Yu, Z., Guo, B., Yang, D., Ma, L., Liu, Z. and Xiong, F., 2022. Data-driven targeted advertising recommendation system for outdoor billboard. ACM Transactions on Intelligent Systems and Technology (TIST), 13(2), pp.1-23.
13. G. Levi and T. Hassner. Age and gender classification using convolutional neural networks. In IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) workshops, June 2015.
14. E. Eidinger, R. Enbar, and T. Hassner. Age and gender estimation of unfiltered faces. Trans. on Inform. Forensics andSecurity, 9(12), 2014S. Baluja and H. A Rowley.Boosting sex identification performance.Int. J. Comput. Vision, 71(1):111–119, 2007.

15. W.-L. Chao, J.-Z.Liu, and J.-J. Ding. Facial age estimation based on label-sensitive learning and age-oriented regression. Pattern Recognition, 46(3):628–641, 2013

16. S. E. Choi, Y. J. Lee, S. J. Lee, K. R. Park, and J. Kim. Age estimation using a hierarchical classifier based on global and local facial features. Pattern Recognition, 44(6):1262–1281,2011.

17. A.C.Gallagher and T. Chen. Understanding images of groups of people. In Proc. Conf. Comput. Vision Pattern Recognition, pages 256–263. IEEE, 2009

18. B. A. Golomb, D. T. Lawrence, and T. J. Sejnowski. Sexnet:A neural network identifies sex from human faces. In NeuralInform. Process. Syst., pages 572–579, 1990d

19. G. Guo, Y. Fu, C. R. Dyer, and T. S. Huang. Image based human age estimation by manifold learning and locally adjusted robust regression. Trans. Image Processing,17(7):1178–1188, 2008.

20. G. Guo, G. Mu, Y. Fu, C. Dyer, and T. Huang. A study on automatic age estimation using a large database. In Proc. Int Conf. Comput. Vision, pages 1986–1991.IEEE, 2009.Neural Inform. Process. Syst., pages 1097–1

21. Y. H. Kwon and N. da Vitoria Lobo.Age classification from facial images. In Proc. Conf. Comput. Vision Pattern Recognition ,pages 762–767. IEEE, 1994