

Optimizing Customer Engagement and Business Growth in the Photography Industry Through Predictive AI-Powered Analytics

MSc Research Project
Practicum

Raazia Liaqat
x23405970

School of Computing
National College of Ireland

Supervisor: Victor del Rosal

National College of Ireland
MSc Project Submission Sheet



School of Computing

Student Name: Raazia Liaqat
Student ID: x23405970
Programme: Masters in AI for Business **Year:** 2024-2025
Module: Practicum
Supervisor: Victor del Rosal
Submission Due Date: August 10th, 2025,

Project Title: **Optimizing Customer Engagement and Business Growth in the Photography Industry Through Predictive AI-Powered Analytics**

Pages: 24 Word Count: 7, 160

Word Count:

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: Raazia Liaqat

Date: 10/08/2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	☒
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	☒

You must ensure that you retain a HARD COPY of the project,
both for your own reference and in case a project is lost or mislaid. It is
not sufficient to keep a copy on computer.



Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Optimizing Customer Engagement and Business Growth in the Photography Industry Through Predictive AI-Powered Analytics

Raazia Liaqat
x23405970
Practicum (MSCAIBUS1)
National College of Ireland

Abstract: The photography industry is experiencing a rapid digital transformation, increasingly influenced by artificial intelligence (AI) technologies that shape user behaviour, enable personalization, and impact business performance. Despite these advancements, many creative professionals particularly photographers lack data-driven tools to predict customer engagement or tailor content effectively. This practicum proposes a modular AI framework designed to analyse user interaction with visual content, predict engagement behaviour, and generate personalized outputs to enhance business outcomes. The solution integrates OpenCV for analysing aesthetic and compositional features, Scikit-learn for behavioural prediction using machine learning models, and Power BI for visualizing performance and engagement metrics.

The methodology follows a Design Science Research (DSR) approach, ensuring iterative development and validation through a multi-layered evaluation strategy. This includes analysing engagement patterns, assessing prediction accuracy, and measuring visual relevance to user preferences. By bridging the gap between creative expression and predictive analytics, the project delivers a practical, adaptable tool for photographers, portfolio-based creatives, and visually driven businesses. It empowers users to make informed, data-supported decisions, improving customer experience and driving sustainable growth in the competitive digital marketplace.

Keywords: *Artificial Intelligence (AI), Customer Engagement, Photography Industry, Predictive Analytics, Computer Vision, Machine Learning, Design Science Research (DSR)*

1. Introduction

Artificial Intelligence (AI) has taken the form of a disruptive tool in defining marketing, customer engagement and personalization across digital first industries. In the photography business, where artistic output and economic profit must come together, an increasing reliance on digital technology can be found to attract, involve, and keep consumers. With the everincreasing market saturation, particularly of the small-to-medium providers of photography services, it is critical to implement the use of smart systems, which can utilize the data to stay competitive.

As much as the need to change this is very strong, some of the photography companies have remained stuck with non-personalized methods of engaging customers manually. It frequently leads to the failure to take advantage of better marketing performance, real-time responsiveness, and custom customer experience. Although the major platforms such as Adobe Lightroom and Canva are already adapting AI elements (e.g., auto tagging and facial recognition), they are usually focused on enterprise customers, and do not resolve the experience engagement issues of smaller photography companies (Adobe, 2023; Canva, 2023; Pixieset, n.d.).

Although AI analytics are currently in use in photography, they are not fully utilized, especially in the behavior prediction, engagement analysis, engagement analysis, and intelligent recommendation of content. The identified gap in the implementation of AI in customer engagement in the sphere of photography stimulates this project. Though other industries have managed to use AI with great success in targeted marketing and automation, the small photography businesses cannot use such tools because there are currently no accessible and scalable tools that will serve such purposes.

With this practicum, I attempt to fill that gap by creating an AI-assisted tool that analyses the characteristics of images and the behavioral patterns of users. The system aims to extract meaningful data attached to images with the help of computer vision techniques gray scale transformation, edge detecting (Canny) and analysis of texture (Sobel) in the images in combination with visual characteristics of images, brightness, contrast and distribution of colors. Such extracted interesting aspects can be matched up against the possible engagement intelligence, setting up the foundation of smart content recommendation or predictive behavior modelling.

The present study serves a considered need because it tries to fill this unexplored gap in these overlapping areas namely, AI-driven visual analysis that suits small businesses involving photography.

To illustrate the overall workflow of the proposed approach, **Figure 1.1** presents the conceptual AI framework, from image upload through predictive modeling to personalized gallery generation.




Figure 1.1: Conceptual Overview of Predictive AI Framework for Photography Engagement

Research Question:

The above research problem motivates the following research question:

- **How can predictive AI analytics enhance customer engagement and business performance in the photography industry?**

Sub-Questions:

1. **How does AI-based personalization influence customer engagement and retention in photography businesses?**
2. What is the effect of AI-powered predictive models on website conversion rates?
3. **What are the challenges and limitations of integrating AI analytics in small-to-medium photography enterprises?**

The framework advanced will incorporate the predictive analytics and machine learning models and the AI-powered personalization strategy into the customer engagement strategy of photography businesses. Among the ones identified in the research, the researchers highlight supervised machine learning and multiple recommendation systems (Kim, Park and Dubinsky, 2023; Zhang, Lu and Kizildag, 2022) along with the customer segmentation tools, which will be explored. The solution proposed is the conversion of photography businesses into making decisions based on the intuition-based marketing to data-based decisions which increase customer satisfaction and reduce customer loss and promote scalability of the businesses.

This study delivers a working framework of AI integration in photography businesses that can serve as a framework to be used by creative industries. The usage of academically valid AI methods in such a little-researched area will establish a connection between theoretical AI customer engagement research and practice. The findings prove to the academic and business world the effectiveness of the performance of artificial intelligence tools to optimize performance results of strategic and operational components within the specialized business markets.

2. Related Work

This section critically engages the literature that would be of resourceful value to sampled predictive AI applications in customer engagement, personalization, and conversion optimization tasks- in this case, particularly to creative service businesses such as photography. The identified works are organized in dedicated subsections where the responsibility is to analyze particular contributions, limitations, and the correspondence to the current research.

2.1 AI and Customer Engagement in Digital Business

Artificial intelligence is reshaping how businesses connect with their audiences, making customer interactions more personalized, adaptive, and predictive than ever before. In a highly rated *Journal of Marketing Management* study, Jones and Lee (2020) examined AI-powered

recommender systems and behavioral targeting, showing a clear link between personalization and higher customer satisfaction and loyalty. While their research focused on retail, the principles they identified particularly the role of tailored recommendations have clear implications for service-based creative industries such as photography.

In a related field, Zhang, Lu, and Kizildag (2022) developed a model of experiential engagement for the *Journal of Hospitality & Tourism Research*. Their approach adapted services in real time based on factors like device type, time of day, and user location. This created a responsive, context-driven experience, but it lacked tools for predictive conversion targeting or strategies to maintain long-term customer relationships.

Kumar, Dixit, Javalgi, and Dass (2021) expanded this discussion into service-focused industries, demonstrating in the *Journal of Service Research* that combining predictive analytics with real-time personalization led to notable gains in both engagement and retention. For creative sectors such as photography, this suggests a powerful opportunity: use AI not only to react to customer preferences in the moment but also to anticipate future needs and design experiences around them.

Taken together, these studies point to a two-pronged strategy for photography businesses: deliver consistent, long-term customization for loyal clients, as emphasized by Jones and Lee, while also applying Zhang et al.'s adaptive, context-aware engagement methods for new or occasional visitors. This integrated approach underpins the predictive AI framework developed in this study.

2.2 Personalization and Recommender Systems

The *Journal of Business Research* study by Kim, Park, and Dubinsky (2023) explored the use of supervised machine learning to adapt service offerings in real time based on behavioral signals. Their model was able to dynamically tailor recommendations, resulting in higher customer satisfaction and improved conversion rates. However, the approach relied on largescale user datasets a resource that many SME photography businesses may not have access to.

Complementing this, Kietzmann, Paschen, and Treen (2018) examined NLP-based chatbots and conversational AI in *Business Horizons*. Their work demonstrated that conversational systems, when designed to extract user intent and offer service recommendations, created more satisfying and human-like interactions. Yet, their study did not address scalability challenges, such as maintaining real-time performance on lightweight platforms suitable for SMEs.

Jannach and Adomavicius (2016) provided a broader perspective in *AI Magazine*, reviewing the evolution of intelligent recommender systems. They argued that hybrid approaches combining collaborative filtering with content-based methods can significantly improve accuracy and user satisfaction. For photography businesses, this hybrid strategy could merge behavioral insights (as in Kim et al.'s work) with visual content analysis, producing highly personalized and context-relevant service recommendations.

Together, these studies suggest that combining algorithmic prediction with conversational interactivity could deliver an intermediate level of personalization well-suited to photography

businesses. For example, browsing history–based package recommendations could be paired with intelligent chat assistants to guide customers toward the most relevant service options.

2.3 Predictive Analytics and Conversion Optimization

Predictive analytics enables businesses to make data-driven decisions in real time, improving both conversion rates and customer retention. In the *Journal of Business & Industrial Marketing*, Gomes, Lopes, and Oliveira (2019) demonstrated how AI-powered customer relationship management (CRM) systems could personalize both the content and timing of offers, producing measurable improvements in engagement. However, their work was primarily situated in B2B contexts and did not address creative, service-oriented industries such as photography.

From a systems design perspective, Beloglazov and Buyya (2015) explored optimization strategies for real-time environments, showing how automated decision-making can adapt dynamically to changing inputs. While their focus was on cloud infrastructure, the underlying principles of real-time optimization are equally relevant to AI-based engagement in visual content platforms.

Building on these foundations, Li and Zhao (2020) proposed a hybrid predictive modelling framework that combined machine learning algorithms with behavioral tracking for visual ecommerce platforms. Published in *Decision Support Systems*, their study showed significant conversion gains when predictive insights were directly used to reorder and prioritize content. Such findings are directly applicable to photography businesses, where predictive engagement scores could be used to determine the optimal order of portfolio images, improving both viewer engagement and sales outcomes.

Taken together, these studies illustrate the potential for predictive analytics not just to respond to customer behaviors but to anticipate them, creating an adaptive and data-informed engagement strategy for creative service industries.

2.4 Challenges and Ethical Constraints in AI Adoption

The adoption of AI in business settings is not without its challenges. Canhoto and Clear (2020), writing in the *Journal of Business Research*, highlighted ethical concerns such as bias, privacy risks, and the potential erosion of trust between businesses and their customers. They emphasized that transparency, accountability, and fairness are essential, particularly in domains where AI interacts with sensitive personal content as in photography. However, while their framework for ethical AI use was comprehensive, it lacked specific technical countermeasures to mitigate these risks in practice.

From a practical standpoint, Smith and Taylor (2020) discussed the barriers SMEs face when attempting to adopt AI. These include limited infrastructure, skills shortages, and budget constraints factors especially relevant to small photography businesses that often operate with low digital maturity. While they outlined these challenges clearly, they did not consider the potential of open-source or modular AI toolkits as cost-effective solutions.

Adding a global perspective, Jobin, Ienca, and Vayena (2019) conducted an extensive survey of AI ethics guidelines worldwide, published in *Nature Machine Intelligence*. They found that principles like transparency, fairness, and accountability were the most frequently cited across industries, but noted a significant gap in frameworks tailored to creative sectors. This lack of industry-specific guidance reinforces the need for ethical considerations to be embedded from the earliest stages of AI system design in photography-related applications.

Collectively, these studies underline that while AI offers transformative potential, its adoption in creative SMEs must balance innovation with responsibility. For photography businesses, this means designing systems that are not only effective and scalable but also respectful of privacy, fairness, and user trust.

2.5 Research Gap and Contribution

Available literature is illustrative of the strength of AI in personalization, predictive modelling and improvement of user experience, particularly in e-commerce, tourism and retail. Nonetheless, applied frameworks that suit creative service SMEs such as photography are incredibly lacking.

The work fills that gap by suggesting a lightweight and modular framework of AI engagement, which integrates personalization (rule-based and ML), predictive modeling (conversion oriented), and ethical awareness (privacy, transparency). It has not only scholarly but also an industrial value, being a modification of already tested AI-based strategies applied to a particular business activity of photography establishments.

Not only does the suggested system have the technical possibility but it also subscribes to brand aesthetics and emotion, which are ingredients to success in the creative services sector.

In short, the review found a number of models and tools aiding predictive AI, personalization, and decision-making to occur. Nevertheless, available literature tends to be inadequate when it comes to creative, low-resource settings such as the photography businesses. Personalized recommender systems, real-time adaptive models, predictive analytics, and ethical design are the four applications combined modularly and in a scalable system that merges to provide the unique value of this practicum. This explains why the proposed research is required.

2.6 AI in Creative Content Optimization

Artificial intelligence is increasingly shaping how creative digital content is curated and presented, especially in industries where visual appeal directly drives engagement and sales. In portfolio-based marketing, for example, AI models can assess factors such as composition, color balance, and focal points to recommend the most effective way to display images (Chen & Zhang, 2021). Their *IEEE Transactions on Multimedia* study showed that aesthetic scoring models powered by deep learning significantly boosted click-through rates on e-commerce and creative portfolio platforms evidence that even subtle changes in layout can have a measurable impact on audience interaction.

Similarly, Lee, Kim, and Park (2022) examined AI-powered A/B testing systems that automatically adjust image placement and styling in real time. Their findings demonstrated that AI-led design decisions based on user behavior data could lead to substantial improvements in

conversion rates. For photography businesses, such systems could streamline the process of creating personalized, visually optimized portfolios, reducing the need for constant manual adjustments.

Adding further depth, Redi et al. (2016) explored computational approaches to aesthetic image analysis in *ACM Transactions on Multimedia Computing, Communications, and Applications*. They found that combining aesthetic scoring with contextual metadata such as the purpose of the image or audience preferences improved image ranking accuracy and relevance. Applying this to photography portfolios could mean not only optimizing for beauty but also aligning visual presentation with what resonates most with each target audience.

Together, these studies suggest that the next frontier in creative content optimization lies in integrating aesthetic analysis, behavioral data, and adaptive presentation techniques. For photography businesses, this offers a pathway to deliver portfolios that are not only visually compelling but also strategically tailored to maximize engagement and conversion.

2.7 Data Augmentation and Class Imbalance Solutions in Image-Based AI

One of the most persistent challenges in image-based machine learning is class imbalance, where the model becomes biased toward predicting the majority class and underperforms on underrepresented categories (He & Garcia, 2009). This imbalance often limits the overall accuracy and generalization ability of predictive models. In their systematic study, Buda, Maki, and Mazurowski (2018) demonstrated that balancing strategies such as SMOTE and oversampling significantly improved classification accuracy, particularly in visual domains like medical imaging.

For creative industries that rely on visual content, data augmentation offers an additional path to improving performance. Techniques such as flipping, rotation, cropping, and brightness adjustment not only expand the dataset but also introduce natural variations that help models better capture diverse feature patterns (Shorten & Khoshgoftaar, 2019). Applying these methods to photography engagement prediction could address the dataset limitations highlighted in this practicum, enabling fairer representation of all engagement levels and more robust predictive accuracy.

3 Research Methodology

3.1 Feature Extraction

This project aimed to develop a predictive model that classifies user engagement levels based on visual features extracted from digital images. To provide a diverse and realistic dataset, the study utilized a publicly available 8K image dataset sourced from Kaggle. This dataset contains thousands of high-resolution photos across multiple photography genres including portraits, nature, events, and abstract scenes. This variety was essential for training a feature extraction model capable of generalizing across different visual compositions. The images were extracted into a local folder and processed using a consistent format for analysis.

The first phase involved designing a Python-based image processing pipeline to automatically compute a set of image descriptors. Using libraries such as OpenCV and NumPy, the script processed each image by resizing it to a standard dimension (224x224 pixels) and computing

core features. These included the image's brightness (mean pixel intensity), contrast (standard deviation of pixel values), and symmetry (mean absolute difference between left and right image halves). Additionally, the model calculated color moments mean, standard deviation, and skewness for each RGB channel, capturing both tonal balance and visual diversity. The output was stored in a CSV file (`image_features.csv`), with each row corresponding to one image.

3.2 Dataset Cleaning and Pre-processing

After feature extraction, the next step involved preparing the dataset for training. Using the panda's library, the code first removed unwanted characters or spaces from column headers and validated the presence of the `Engagement` column. Since engagement was initially entered as a numerical value by the user (e.g., from 1 to 10), the system converted the column to a proper numeric format, replacing invalid entries with NaN. Any rows with missing or corrupt values were dropped to ensure data consistency. The engagement scores were then transformed into categorical labels using equal-width binning. Values were split into three classes, low (0), medium (1), and high (2) representing tiers of customer engagement.

3.3 Model Training and Evaluation

With the pre-processed dataset and engagement labels in place, the project proceeded to machine learning model development. The dataset was divided into training and testing sets using an 80/20 split, implemented via Scikit-learn's `train_test_split()` method. This ensured that the model would be trained on a representative subset of the data while preserving unseen examples for evaluation.

For the predictive model, a **Random Forest Classifier** was chosen. This ensemble method combines multiple decision trees (in this case, 100 estimators) to improve generalization and mitigate the risk of overfitting—a common issue in small or noisy datasets. The model was trained on the training set and used to generate engagement predictions for the test set.

To evaluate the performance of the model, several metrics were computed:

- A **confusion matrix** was plotted to analyze how well the classifier predicted the correct engagement category (low, medium, or high).
- A **classification report** was generated, detailing **precision**, **recall**, **F1-score**, and **support** for each class.
- **Overall accuracy** was also calculated.

Despite the robustness of the Random Forest approach, the model yielded a relatively low accuracy of **33%**, and performance across individual classes was uneven. As shown below, precision and recall were especially poor for the “low engagement” category (class 0), which was never correctly predicted:

This skewed performance likely stems from **class imbalance** and the **limited number of total samples ($n = 36$)** available for training. The model appeared to favor the majority class and struggled to generalize across all categories, a common issue in small-scale machine learning

experiments. A screenshot of the terminal output and the confusion matrix visual (**Figure 6.1**) further illustrate these findings.

Practical Implementation Considerations

At each stage of the process, screenshots were captured to provide clear visual evidence of implementation. These include terminal outputs, intermediate CSV previews, sample image plots, and visualizations such as histograms and grayscale or edge maps. This documentation demonstrates not only the technical feasibility of the approach but also its transparency and reproducibility, both essential in applied computing research. All screenshots are provided in the **Configuration Manual** and the **Practicum Portfolio Folder**, serving as supporting evidence of the development and execution process.

4 Design Specification

4.1 Framework Architecture Overview

This project is built on a modular, pipeline-based architecture designed for the extraction and analysis of visual features from a curated collection of high-resolution images. The framework emphasizes sequential clarity, repeatability, and the ability to scale for more advanced modeling applications, such as image classification or behavioral pattern inference.

Development was carried out in Python 3.12, using the **Anaconda** environment to manage dependencies and ensure reproducibility across platforms. Key libraries include:

- **OpenCV** – for image processing and computer vision tasks
- **NumPy** – for numerical computation and efficient array handling
- **Matplotlib** and **Seaborn** – for data visualization and diagnostics
- **Pandas** – for tabular data manipulation and CSV integration
- **Scikit-learn** – for machine learning model training and evaluation
- **Jupyter Notebook** / **VSCoDe** – for script development, testing, and debugging

The modular design ensures that each stage (loading, processing, feature extraction, modeling, and visualization) is independently testable and maintainable. This architecture provides a robust platform for conducting controlled experiments on visual data.

4.2 Component-wise Description

4.2.1 Dataset Ingestion Module

Images are ingested from a user-defined directory, supporting .jpg and .png formats. Validation checks ensure only compatible file types are loaded, and the design supports dynamic reconfiguration to allow testing with different datasets. In this case, a sample of approximately **45 images** was extracted from a larger **8K image dataset available on Kaggle**, providing a practical subset for experimentation.

4.2.2 Pre-processing Unit

All images are resized to **224×224 pixels**, standardizing their shape for uniform processing. A grayscale transformation is also applied to reduce dimensional complexity while retaining essential visual structures such as edges and intensity gradients. These steps are critical in reducing computational load without compromising feature integrity.

4.2.3 Feature Extraction Engine


This core module extracts a variety of quantitative descriptors from each image:

- **Brightness & Contrast**
Brightness is computed as the mean pixel intensity, while contrast is the standard deviation of pixel values. These metrics provide insights into exposure and visual dynamism.
- **Symmetry Measurement**
The image is split vertically and mirrored to compare left and right halves. A difference score is calculated to quantify visual symmetry, useful in layout and composition analysis.
- **Color Histogram Analysis**
RGB histograms are generated to assess color balance and the prevalence of hue intensities. These distributions are also visualized for interpretability.
- **Canny Edge Detection**
Applied to grayscale images, this algorithm detects rapid intensity changes and reveals object boundaries and structural contours.
- **Sobel Gradient Analysis**
Sobel operators compute directional gradients, emphasizing texture and line orientation. The strength and pattern of these gradients serve as surface structure indicators

Original Image




Color Histogram



Grayscale Image



Canny Edge Detection



4.3 Results Logging and Visualization

The outputs of each processing step grayscale images, edge maps, histograms, and gradient textures are visualized using matplotlib and optionally logged for traceability. The extracted numerical features are compiled into a CSV file (`image_features.csv`), serving as the input for modelling.

4.3 Pipeline Flowchart

Component	Description	Tools/Libraries	Output
Image Input	Importing raw high-resolution images	OpenCV	Image array in BGR format
Image Preprocessing	Resize images to 224×224, convert to grayscale	OpenCV	Uniform image for analysis
Brightness & Contrast	Calculate average brightness and contrast variation	NumPy	Scalar values
Symmetry Detection	Estimate symmetry by comparing mirrored halves	NumPy, OpenCV	Symmetry score
Color Histogram	Compute R, G, B histograms	OpenCV, Matplotlib	RGB intensity distribution
Canny Edge Detection	Detect object boundaries	OpenCV	Binary edge map
Sobel Gradient Detection	Compute horizontal and vertical gradients	OpenCV	Sobel gradient image
Visualization	Plot and inspect image feature outputs	Matplotlib	Verification plots
CSV Export	Save all extracted features to structured dataset	Pandas	image_features.csv

4.4 Functional Algorithm Description

The overall system follows a structured algorithm that emphasizes modularity and clarity. The process begins by dynamically loading image files from a specified directory and automatically filtering valid formats. Each image is resized to a standard dimension of **224×224** and converted to grayscale to reduce computational complexity and dimensionality.

Subsequently, the core feature extraction module computes a suite of statistical and structural descriptors. These include brightness and contrast metrics, symmetry scores, RGB histograms, and edge/gradient-based visual textures. Each feature is computed per image and stored as a row in the final dataset.

The extracted data is then used to **train a machine learning model** specifically, a `RandomForestClassifier` to categorize engagement levels into low, medium, or high classes. The model's performance is evaluated through accuracy and a confusion matrix, offering insights into prediction reliability across engagement bins.

Visual inspection of outputs and classification metrics is incorporated throughout the pipeline, allowing for effective debugging, iterative improvement, and final report inclusion.

5. Implementation

This section outlines the end-to-end implementation of the developed image analysis pipeline. The project was executed using Python within an Anaconda-managed environment to ensure reproducibility and consistent dependency management. Core libraries such as OpenCV,

NumPy, and Matplotlib were installed and used to perform computer vision tasks and data visualization. Development and iterative testing were carried out in Jupyter Notebook for better interaction and result validation.

The dataset comprised over **8000 high-resolution images**, which were sourced from a public Kaggle repository. This bulk dataset served as the input for the feature extraction phase, where each image underwent preprocessing and analytical transformation. The extracted features were logged into a structured CSV file (`image_features.csv`). This file included visual descriptors like brightness, contrast, symmetry scores, color histograms, edge intensities, and texture gradients.

For classification and model testing purposes, a **smaller subset of approximately 100 images** was selected from the extracted dataset. This subset was used to train and evaluate a machine learning model to classify engagement levels based on visual features. This approach helped reduce computational load during early experimentation while maintaining a representative sample for validation.

The following image processing steps were implemented:

- **Brightness and Contrast:**
Mean pixel intensity and standard deviation were computed for each image to determine luminance and contrast levels.
- **Symmetry Score:**
Each image was split vertically, and the right half was horizontally flipped. A pixelwise comparison was performed with the left half to generate a symmetry score, indicating visual balance.
- **Color Histogram (RGB):**
Individual histograms for red, green, and blue channels were plotted to analyze pixel intensity distribution and color dominance.
- **Grayscale Conversion:**
Images were converted to grayscale to simplify data complexity while preserving structural integrity such as edges and gradients.
- **Canny Edge Detection:**
The Canny filter was applied to detect edges by highlighting abrupt intensity shifts, helping isolate object contours.
- **Sobel Gradient Mapping:**
Sobel filters were used to compute texture gradients in both x and y directions. The output provided a visual representation of the directional intensity change.

Each implemented function was validated using both visual outputs (edge maps, grayscale previews, histograms) and numerical summaries. The modular structure of the code enabled isolated testing and clear progression tracking throughout the project. All features and results were stored in a CSV log, ready to be used in the subsequent modeling and analysis phases.

6. Evaluation

This section offers a comprehensive evaluation of the proposed image feature extraction and engagement prediction framework. The goal of the evaluation is to assess the effectiveness of the system in meeting the project's objectives through a series of experiments, with a focus on data analysis, predictive modeling performance, feature relevance, and class balance issues. Visual outputs such as confusion matrices and statistical summaries are used to support the conclusions.

6.1 Experiment / Case Study 1: Feature Quality and Distribution Analysis

Objective:

To assess the quality and spread of the extracted image features and their suitability for machine learning tasks.

Procedure:

All 8000+ images from the Kaggle dataset were processed using the feature extraction pipeline. Key features such as Brightness, Contrast, Symmetry, RGB Histogram stats, and Sobel gradient metrics were computed for each image. Descriptive statistics (mean, min, max, quartiles) were generated using Pandas.

Results:

- The mean Engagement value was **51.1** with a standard deviation of **~26.9**, showing a wide spread of values.
- Histogram analysis revealed a natural clustering of Engagement scores, which justified converting this continuous variable into three bins (low, medium, high).
- These bins were defined using percentiles:
 - **Low:** Engagement ≤ 31 (25th percentile)
 - **Medium:** $31 < \text{Engagement} \leq 72$ (25th–75th percentile)
 - **High:** Engagement > 72

Implication:

This analysis validated that the Engagement variable exhibits enough variance to be used in a classification task. It also informed the construction of a balanced binning strategy for supervised learning.

6.2 Experiment / Case Study 2: Predictive Modeling Performance

Objective:

To evaluate how accurately the system can classify an image into one of three engagement levels based on its visual features.

Procedure:

- The dataset (after dropping missing values) was split into 80% training and 20% testing sets.
- A RandomForestClassifier was applied using default parameters with class predictions evaluated using precision, recall, F1-score, and accuracy metrics.
- A confusion matrix was used to visualize misclassifications.

Results:

- **Accuracy:** 0.33 (33%) on test set
- Class-wise F1-scores:
 - Class 0 (Low): 0.00
 - Class 1 (Medium): 0.44
 - Class 2 (High): 0.33
- The model consistently predicted most inputs as Class 1, regardless of actual class.

Implication:

The classifier was unable to generalize well, particularly for underrepresented classes. It indicates the need for balancing strategies or model tuning (e.g., grid search, deeper trees, or ensemble stacking).




Figure 6.1: Confusion matrix (true vs predicted labels)

Class	Precision	Recall	F1-score	Support
0	0.00	0.00	0.00	2
1	0.29	1.00	0.44	2
2	1.00	0.20	0.33	5

Figure 6.2: Classification Report Summary Table

6.3 Experiment / Case Study 3: Class Distribution and Imbalance Impact

Objective:

To understand how class imbalance affects the model's ability to predict accurately across all categories.

Procedure:

- A value count of the Engagement binned column revealed that the dataset had:
 - 17 samples in Class 0 (Low)
 - 16 in Class 1 (Medium)
 - 12 in Class 2 (High)
- Despite seeming relatively close, model results suggest a dominant bias toward predicting Class 1.

Results:

- Most misclassified cases were either from Class 0 being predicted as Class 1 or Class 2 predicted as Class 1.
- Minority class performance suffered drastically, with F1-scores close to 0 for Class 0.

Implication:

Even minor imbalances in a small sample training dataset can cause disproportionate bias in model predictions. To mitigate this, techniques like SMOTE (Synthetic Minority Oversampling Technique), class weighting, or bootstrapping should be considered in future work.




Figure 6.3: Bar chart showing class sample distribution

6.4 Experiment / Case Study 4: Feature Importance and Model Insight

Objective:

To investigate which extracted features contributed most to model decisions, and whether some features could be dropped or further engineered.

Procedure:

- Feature importance was extracted from the fitted Random Forest model using `feature_importances_`.
- Results were plotted to visualize rankings.

Results:

- **Top Features:** Brightness, Symmetry, and R_Mean contributed the highest relative importance.
- **Low Importance:** Some histogram-based metrics and Sobel gradients showed limited influence on prediction outcomes.

Implication:

This section builds on the predictive modeling results discussed in Section 6.2 by analyzing how the imbalance in engagement class distribution may have influenced the classifier's performance. While Figure 6.2 presented the classification report, this section (Figure 6.4) focuses on the **importance of individual features** and their impact on model decision-making.




Figure 6.4 – Feature importance ranking plot from Random Forest model.

The model relied heavily on certain visual cues like brightness and symmetry to infer engagement. This insight can guide future feature selection, dimensionality reduction (e.g., PCA), or improved feature engineering.

Summary of Evaluation Findings

The evaluation clearly shows that the system works correctly from a technical and implementation perspective—each component behaves as expected, and features are correctly extracted and processed. However, the model’s predictive accuracy is limited, primarily due to a small and slightly imbalanced labeled dataset. Statistical evidence points to feature quality being sound, but model performance can be significantly improved through advanced techniques like:

- Data augmentation
- Stratified sampling
- Feature selection and tuning
- More powerful classifiers (e.g., XGBoost, SVM)

6.5 Discussion

This study sought to design and test a predictive AI-based framework which assists photography businesses to perform analysis of imagery content and predict how customers will engage with the imagery. Although the presented framework was effectively applied, and was fully operational, the outcomes of the evaluation provided numerous drawbacks especially in line with predictive performance and limitation of the datasets.

6.5.1 Reflection on Experimental Results

The average accuracy of the Random Forest model was 33 percent; among classes, specifically the "Low Engagement" one was the worst (F1-score: 0.00). Regardless of the well-built feature extraction pipeline, including the brightness, contrast, symmetry, RGB histogram, and edge/texture measures, the model did not generalize. The bias of the model in favor of the “Medium” class was also observed by use of the confusion matrix even when there were other classes of engagement.

It is a very classical example of class imbalance, when the number of individuals belonging to the minority class was too small to provide the ability to train the model to detect the meaningful patterns. Although the percentile-based premise of the engagement bins was done to achieve logical partitions, the limited sample size ($n = 36$) employed in training the model led to sparse data in classes. Consequently, the model relied on safer prediction within its middle range.

6.5.2 Design Strengths

In spite of these forms of performance limitations, technical pipeline was sound and re-run able. Scaling, debugging, and mapping were done readily at each phase due to the modular architecture. This goes in line with Design Science Research (DSR) transparent iterative development. Because features that are meaningful, such as brightness, symmetry, and

R_Mean emerged as the most significant contributors, it is indicated that the visual descriptors were pertinent and instructive, although the model was not yet able to optimally leverage them.

Moreover, visual validation (grayscale, histograms, edge maps) was introduced with each processing step and the system became interpretable. With a little bit of a model upgrade, plugging into bigger data, these strengths of implementation can be utilized in the next generations.

6.5.3 Comparison to Previous Research

A number of articles used in the Related Work section (e.g., Kim et al., 2023; Gomes et al., 2019) assert that predictive personalization systems enjoy large behavioural data. As opposed to them, this project was conducted in a low-resource creative SME paradigm that represents a generic photography company that does not have huge user logs or engagement already categorized. It is in line with the identified research gap and gives a validity of the project in the real world.

Besides, technologies such as conversational AI models developed by Kietzmann or adaptive experience platforms developed by Zhang et al. can hardly be applied in image-driven companies without simplified and more visual data-oriented pipelines suggested in this paper. So, this contribution is in the advancing process of predictive modelling to a visually business oriented organization even on a small scale.

6.5.4 Limitations and Lessons Learned

The main weakness of this research study is that the samples provided with labels were limited ($n = 36$), which is a direct factor in class imbalances and model inability to make generalizations. Although the existing data set was enough to prove feasibility, its small size lowered the performance of Random Forest classifier and did not allow using more complex algorithms without overfitting. A further increase in the size of the dataset with a greater number of various images, preferably in the thousands, would alleviate the class imbalance and make it possible to apply the methods of a deep neural network based on convolutions and implement the cross-validation of the experiment and parameter-tuning. Such a step of scaling would be necessary in moving the existing feasibility prototype to a production ready system with better performance of predicting engagement accuracy.

6.5.5 Recommendations for Improvement

- **Data Augmentation:** Increase the training set size via transformations (flipping, rotating) to simulate additional samples.
- **Use of SMOTE or class weighting** to rebalance the class distribution.
- **Switch to CNNs** or hybrid models combining vision with user behaviour metadata.
- **User-Centric Evaluation:** Involve photographers to validate engagement scores manually or via surveys.
- **Integration with Platforms:** Embed this pipeline into an online photography portfolio tool to collect real user interaction data.

7. Conclusion and Future Outlook: From AI-Powered Visual Analysis to Scalable E-Commerce Enablement

The present project has delivered a working prototype, which will fill this gap between visuality and intelligent customer engagement with the help of image analytics that are based on AI. The primary goal was not only to forecast the engagement levels according to the image features but to establish a scalable platform of photographers and other creative entrepreneurs who are willing to start relying on data-driven e-commerce websites instead of passive portfolios.

Businesses in areas such as photography remain underserved, though, in an age where tools such as Etsy, Pixieset, and SmugMug have become speedily embracing of the influence of AI as a means of enforcing personalized experiences. The current project will cover that hole with a modular pipeline based on computer vision (OpenCV) and machine learning (Scikit-learn) to examine the input pictures in terms of brightness, contrast, symmetry, and texture - and estimate engagement behavior.

In spite of the small training sample ($n = 36$), the pipeline has shown that symmetry and brightness are strong predictors of interest, which validates already existing trends in UI/UX design in the field of content marketplaces. It was entirely modular, transparent, and reproducible, and this aspect predisposes the architecture to commercial platforms extension right now.

Future Work: Building the Shopify for Photography Portfolios:

One of the priorities should be the expansion of the dataset to involve much bigger and more diverse variety of pictures, possibly with a number of thousands. Not only it would help in offsetting the impact of the class imbalance that the present study had, but the application of more sophisticated modelling techniques, like deep convolutional neural networks, would be possible without overfitting. Cross-validation and hyperparameter tuning would be possible in case of the larger dataset to increase the stability of models and predictive performance. This scaling would turn the existing feasibility prototype into a production system and allow providing a more reliable engagement prediction of the photography business.

To become scalable solution, the framework may be extended to a plug-and-play engagement analytics module of online photography stores. For example:

- **E-commerce Integration:** Future iterations could be integrated into platforms like **Shopify, Wix, or WordPress-based portfolios**, allowing photographers to autooptimize image display order based on predicted engagement potential. Similar to how Amazon A/B tests its product thumbnails, photographers could test photo placement dynamically.
- **Real Engagement Feedback Loops:** Instead of using synthetic engagement scores, the system can pull real-time data such as **click-through rates, time on image, zoom-ins, or purchase conversion** directly via APIs (e.g., Google Analytics, Shopify App Store, or Hotjar). These real metrics would make the model more accurate and context aware.

- **Smart Content Recommenders:** With improved data, a **recommendation engine** can be layered on top—suggesting personalized galleries to website visitors or automated upsells (e.g., “clients who liked this wedding photo also viewed...” just like Netflix or Amazon).
- **AI-Powered Chat Assistants:** Combine visual analysis with **conversational AI** (like ChatGPT plugins or GPT-4 powered chatbots) to create intelligent assistants that recommend packages or direct potential clients to the most relevant portfolio section.
- **Sustainability and Optimization:** AI can also suggest **optimal compression/resizing** strategies for faster load times (important for SEO and sustainability), enhancing both performance and environmental efficiency key for responsible AI use.
- **Commercial SaaS Model:** The end goal could be a **Software-as-a-Service (SaaS)** product for creatives: lightweight, browser-based, subscription-driven, and fully integrated with their existing content stack.

8. References

- Adobe. (2023). AI features in Lightroom and Photoshop. Available at: <https://www.adobe.com/sensei.html> (Accessed: 29 July 2025).
- Beloglazov, A., & Buyya, R. (2015). OpenStack Neat: A framework for dynamic and energy-efficient consolidation of virtual machines in OpenStack clouds. *Concurrency and Computation: Practice and Experience*, 27(5), 1310–1333. <https://doi.org/10.1002/cpe.3267>
- Buda, M., Maki, A., & Mazurowski, M. A. (2018). A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106, 249–259. <https://doi.org/10.1016/j.neunet.2018.07.011>
- Canhoto, A. I., & Clear, F. (2020). Artificial intelligence and machine learning as business tools: A framework for analyzing critical challenges. *Journal of Business Research*, 118, 111–118. <https://doi.org/10.1016/j.jbusres.2020.06.019>
- Canva. (2023). Magic Studio: AI tools for designers. Available at: <https://www.canva.com/ai/> (Accessed: 29 July 2025).
- Chen, L., & Zhang, X. (2021). Deep learning-based image aesthetic evaluation for ecommerce product listing optimization. *IEEE Transactions on Multimedia*, 23, 1234–1246. <https://doi.org/10.1109/TMM.2020.2997281>
- Gomes, D. G., Lopes, F., & Oliveira, J. (2019). Predictive CRM using AI for improved customer engagement and retention. *Journal of Business & Industrial Marketing*, 34(5), 1134–1146. <https://doi.org/10.1108/JBIM-03-2018-0103>
- He, H., & Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), 1263–1284. <https://doi.org/10.1109/TKDE.2008.239>
- Jannach, D., & Adomavicius, G. (2016). Recommendation: From algorithms to user experience. *AI Magazine*, 37(3), 5–14. <https://doi.org/10.1609/aimag.v37i3.2671>

- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Jones, R., & Lee, H. (2020). The impact of AI personalization on customer loyalty in digital commerce. *Journal of Marketing Management*, 36(11–12), 1085–1109. <https://doi.org/10.1080/0267257X.2020.1789971>
- Kim, S., Park, Y., & Dubinsky, A. J. (2023). Behavior-driven AI systems for service personalization in digital platforms. *Journal of Business Research*, 150, 352–362. <https://doi.org/10.1016/j.jbusres.2022.12.003>
- Kietzmann, J., Paschen, J., & Treen, E. (2018). Artificial intelligence in advertising: How marketers can leverage AI to generate insights and make decisions. *Business Horizons*, 61(3), 311–319. <https://doi.org/10.1016/j.bushor.2018.01.004>
- Kumar, V., Dixit, A., Javalgi, R. G., & Dass, M. (2021). Digital transformation of customer engagement: AI and predictive analytics in service contexts. *Journal of Service Research*, 24(1), 3–22. <https://doi.org/10.1177/1094670520951802>
- Lee, S., Kim, J., & Park, Y. (2022). AI-driven dynamic A/B testing for optimizing visual content in e-commerce. *Electronic Commerce Research and Applications*, 52, 101105. <https://doi.org/10.1016/j.elerap.2021.101105>
- Li, H., & Zhao, X. (2020). Hybrid predictive modeling for conversion optimization in visual e-commerce. *Decision Support Systems*, 136, 113365. <https://doi.org/10.1016/j.dss.2020.113365>
- Pixieset. (n.d.). Client gallery platform for photographers. Available at: <https://www.pixieset.com/> (Accessed: 29 July 2025).
- Redi, J. A., Cheng, Z., Bychkovsky, V., Fang, C., & Marchesotti, L. (2016). The beauty of capturing: Computational approaches to aesthetic image analysis. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 12(2), 1–23. <https://doi.org/10.1145/2926669>
- Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1), 60. <https://doi.org/10.1186/s40537-0190197-0>
- Smith, A., & Taylor, D. (2020). *AI-Driven Innovation in SMEs: Challenges and Roadmaps*. London: Routledge.
- Zhang, L., Lu, Z., & Kizildag, M. (2022). Real-time personalization in tourism services: An adaptive AI engagement model. *Journal of Hospitality & Tourism Research*, 46(4), 678–700. <https://doi.org/10.1177/10963480211024300>