

Advancing Machine Learning: Development, Evaluation, and Feature Engineering in Domain-Specific Applications

Fareesa Khan¹

Lecturer, Department of Electronic Engineering
QUEST, Larkana, Pakistan
enr.fareesa.khan@quest.edu.pk

Muhammad Ibrahim Channa²

Professor, Department of Information Technology,
Quaid-e-Awam University of Engineering, Science and Technology, Nawab Shah, Pakistan
Ibrahim.channa@quest.edu.pk

Mohammad Ali Soomro³

Assistant Professor Department of Computer Engineering,
Quaid-e-Awam University of Engineering, Science and Technology, Nawab Shah, Pakistan
kalakot2@gmail.com

Shah Zaman Nizamani⁴

Assistant Professor, Department of Information Technology,
Quaid-e-Awam University of Engineering, Science and Technology, Nawab Shah, Pakistan
shahzaman@quest.edu.pk

Muhammad Aamir Bhutto⁵

Assistant Professor, Department of Software Engineering,
Quaid-e-Awam University of Engineering, Science and Technology, Nawab Shah, Pakistan
engraamirbhutto@gmail.com

Abstract— The rapid advancements in machine learning and the increasing availability of extensive datasets have significantly propelled the field of image classification. This study presents a comprehensive evaluation of three prominent machine learning models—Convolutional Neural Networks (CNNs), k-Nearest Neighbors (kNN), and Random Forest classifiers—on a specific image classification task. The research investigates the efficacy of these models through various performance metrics, examining their strengths and limitations. CNNs demonstrated superior accuracy and robustness, attributed to their ability to learn hierarchical features directly from image data. However, they require substantial computational resources and large datasets. The kNN classifier, while straightforward and easy to implement, exhibited limitations in handling high-dimensional data. The Random Forest classifier showed promise in structured data analysis but required effective feature engineering to enhance its performance with image data. The study also highlights the critical role of feature engineering techniques, data preprocessing, and hyperparameter tuning in optimizing model performance. Advanced CNN architectures, ensemble methods, and real-time deployment strategies are proposed as future research directions to further enhance image classification systems. This research provides valuable insights for developing more accurate and efficient image classification models, with potential applications across various domains.

Keywords- *Convolutional Neural Networks, k-Nearest Neighbors, Random Forest, Image Classification, Feature Engineering*

I. INTRODUCTION

In recent years, the field of machine learning has experienced significant advancements, driven by the exponential growth in data availability and the increasing computational power of modern hardware. Among the various domains benefiting from these advancements, image classification has emerged as a pivotal area of research and application. From medical diagnosis to autonomous driving and facial recognition, image classification technologies are transforming numerous sectors by automating and enhancing the accuracy of visual data interpretation [1]. Convolutional Neural Networks (CNNs), k-Nearest Neighbors (kNN), and Random Forest classifiers are among the most widely studied and utilized algorithms in this domain, each bringing unique strengths and capabilities to the task of image classification [2]. CNNs have revolutionized image classification with their ability to automatically learn hierarchical feature representations from raw pixel data. Unlike traditional methods that rely on handcrafted features, CNNs can capture complex patterns through multiple layers of convolutional filters, pooling operations, and non-linear activations [3]. This deep learning approach has led to state-of-the-art performance in various image recognition tasks, outperforming previous methods in terms of accuracy and generalization. However, CNNs require large amounts of labeled data and significant computational resources, which can be a barrier to their widespread adoption.

On the other hand, kNN is a simpler, instance-based learning algorithm that classifies a new sample based on the majority class among its nearest neighbors in the feature space. Despite its simplicity and ease of implementation, kNN can suffer from high computational costs and poor performance with high-dimensional data [4]. Random Forest classifiers, which are ensemble methods based on decision trees, offer another alternative by combining the predictions of multiple trees to improve robustness and accuracy. They are particularly effective in handling structured data and can provide insights into feature importance, but may still face challenges with image data due to their reliance on pre-engineered features. Despite the progress made in image classification, several challenges persist. One significant challenge is achieving high accuracy and robustness across diverse and complex datasets. Traditional methods like kNN and Random Forest often struggle with high-dimensional data and intricate patterns, while even sophisticated CNN architectures can be prone to overfitting and require substantial computational resources [5]. Furthermore, the effectiveness of these models is heavily dependent on the quality of feature engineering and data preprocessing steps. This paper aims to address these challenges by systematically evaluating the performance of CNNs, kNN, and Random Forest classifiers on a specific image classification task. We will also investigate the impact of various feature engineering techniques and data

preprocessing methods on the overall performance of these models.

Through this comprehensive analysis, we seek to identify the strengths and limitations of each approach and provide insights into the most effective strategies for improving image classification accuracy and reliability. By comparing the performance of CNNs, kNN, and Random Forest classifiers, we aim to offer a nuanced understanding of how these models can be optimized for different types of image data. Additionally, we will explore advanced techniques such as hyperparameter tuning, data augmentation, and the use of ensemble methods to further enhance model performance. This research will contribute to the ongoing development of more accurate, efficient, and interpretable image classification systems, with potential applications across a wide range of fields.

II. LITERATURE REVIEW

Machine learning (ML) has revolutionized various industries by providing advanced tools and methodologies to solve complex problems across numerous domains. In domain-specific applications, ML is tailored to address the unique challenges and requirements inherent to particular fields such as healthcare, finance, agriculture, and more. The primary advantage of ML in these applications is its ability to learn from vast amounts of data, identify patterns, and make informed predictions or decisions without explicit programming for every task. In healthcare, for example, ML algorithms have been extensively utilized to improve diagnostic accuracy, predict patient outcomes, and personalize treatment plans [6]. Techniques such as convolutional neural networks (CNNs) have shown remarkable success in medical image analysis, aiding in the detection of diseases such as cancer, diabetic retinopathy, and cardiovascular conditions. The ability of ML models to process and analyze complex medical images with high precision has led to significant advancements in early disease detection and treatment efficacy.

The finance sector also benefits greatly from ML applications. Algorithms are used for credit scoring, fraud detection, algorithmic trading, and risk management. By analyzing historical data and identifying trends, ML models can predict market movements, assess creditworthiness, and detect fraudulent transactions with higher accuracy than traditional methods [7]. The integration of ML into financial systems enhances operational efficiency, reduces risks, and provides better customer service through personalized financial advice and products. Agriculture is another domain where ML has made substantial contributions. Precision agriculture leverages ML techniques to optimize farming practices, enhance crop yields, and manage resources efficiently. Sensors and drones collect data on soil conditions, weather patterns, and crop health, which is then analyzed by ML models to provide actionable insights for farmers [8]. This data-driven approach helps in making informed

decisions about irrigation, fertilization, and pest control, leading to sustainable farming practices and increased productivity.

In addition to these examples, ML is applied in various other fields such as transportation, manufacturing, and marketing. In transportation, ML algorithms improve route optimization, traffic management, and autonomous vehicle navigation. Manufacturing industries use ML for predictive maintenance, quality control, and supply chain optimization. In marketing, ML helps in customer segmentation, sentiment analysis, and targeted advertising, enabling businesses to understand and cater to customer preferences more effectively [9]. The versatility of ML in domain-specific applications stems from its ability to adapt to the unique data characteristics and requirements of each field. By leveraging domain-specific knowledge, feature engineering, and customized algorithms, ML models can deliver highly accurate and relevant solutions. The continuous advancements in ML research, including the development of more sophisticated algorithms and the availability of large datasets, further enhance its potential to address complex challenges in various domains [10].

A. Feature Engineering Techniques

Feature engineering is a critical step in the machine learning pipeline, involving the creation, transformation, and selection of relevant data features to improve model performance [11]. Effective feature engineering can significantly enhance a model's predictive power and accuracy by providing it with the most informative and discriminative features. Several techniques are commonly used in feature engineering, each tailored to the nature of the data and the specific requirements of the application [12]. One fundamental technique is scaling and normalization, which adjusts the range of features to a standard scale. This is particularly important for algorithms sensitive to the scale of data, such as support vector machines (SVMs) and k-nearest neighbors (KNN). Common methods include Min-Max Scaling, which rescales features to a fixed range (typically 0 to 1), and Z-Score Normalization, which transforms features to have a mean of zero and a standard deviation of one.

Feature transformation involves creating new features from existing ones through mathematical operations [13]. Techniques such as logarithmic scaling, polynomial features, and interactions between features can help capture non-linear relationships and interactions in the data. For instance, polynomial features create new features by taking polynomial combinations of the original features, which can help linear models capture non-linear relationships. Encoding categorical variables is another crucial aspect of feature engineering, especially when dealing with non-numeric data. Techniques such as one-hot encoding, label encoding, and target encoding are used to convert categorical data into a numeric format that machine learning algorithms can process. One-hot encoding creates binary columns for each category, while label encoding

assigns a unique integer to each category [14]. Target encoding, on the other hand, replaces categories with the mean target value for that category, which can be particularly useful for high-cardinality categorical features.

Dimensionality reduction techniques, such as Principal Component Analysis (PCA) and t-distributed Stochastic Neighbor Embedding (t-SNE), are used to reduce the number of features while preserving the most important information. PCA transforms the original features into a set of linearly uncorrelated components, ordered by the amount of variance they capture from the data. t-SNE, primarily used for visualization, reduces dimensions by modeling the probability distribution of data points in high-dimensional space and projecting them into a lower-dimensional space [15]. Handling missing values is another critical aspect of feature engineering. Techniques like imputation, where missing values are replaced with statistical measures (mean, median, mode) or more sophisticated methods such as KNN or regression-based imputation, are commonly used. Additionally, creating new features to indicate the presence of missing values can also be helpful, as the pattern of missingness itself might carry significant information.

Feature selection is the process of identifying the most relevant features for the model. Techniques such as recursive feature elimination (RFE), LASSO (Least Absolute Shrinkage and Selection Operator), and tree-based feature importance rankings help in selecting features that contribute the most to the predictive power of the model. By removing irrelevant or redundant features, feature selection can improve model interpretability and reduce overfitting.

B. Evaluation Metrics in Machine Learning

Evaluating machine learning models accurately is crucial for understanding their performance and effectiveness. Different metrics are used depending on the type of problem (classification, regression, clustering, etc.) and the specific goals of the model. Proper evaluation helps in selecting the best model, tuning its parameters, and ensuring its reliability and robustness [16].

In classification problems, several metrics are commonly used. Accuracy measures the proportion of correctly classified instances out of the total instances but can be misleading in imbalanced datasets [17]. Precision and recall provide a more detailed understanding of model performance. Precision indicates the proportion of true positive predictions among all positive predictions, while recall (or sensitivity) measures the proportion of true positive predictions among all actual positives. The F1 score combines precision and recall into a single metric, offering a balance between the two, especially useful when the classes are imbalanced. Receiver Operating Characteristic (ROC) curves and the Area Under the Curve (AUC) provide a graphical representation of a classifier's

performance across different threshold values, highlighting its ability to distinguish between classes [18].

For regression problems, metrics such as Mean Absolute Error (MAE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE) are commonly used. MAE measures the average magnitude of errors in a set of predictions, without considering their direction. MSE, on the other hand, squares the errors before averaging them, penalizing larger errors more severely. RMSE is the square root of MSE, providing a measure of the average error magnitude in the same units as the original data. R-squared (coefficient of determination) is another key metric that indicates the proportion of variance in the dependent variable that is predictable from the independent variables, offering insight into the model's explanatory power. In clustering problems, metrics such as Silhouette Score, Davies-Bouldin Index, and Adjusted Rand Index (ARI) are often used. The Silhouette Score measures how similar an object is to its own cluster compared to other clusters, indicating the coherence of clusters [19]. The Davies-Bouldin Index evaluates the average similarity ratio of each cluster with its most similar cluster, with lower values indicating better clustering. The Adjusted Rand Index measures the similarity between the predicted clustering and a ground truth clustering, adjusting for the chance grouping of elements.

Cross-validation techniques are essential for robust model evaluation. K-fold cross-validation involves partitioning the dataset into K subsets and training the model K times, each time using a different subset as the validation set and the remaining subsets as the training set [20]. This method helps in assessing the model's performance more reliably by reducing the bias associated with a single train-test split. Stratified K-fold cross-validation is a variation where each fold maintains the same class distribution as the entire dataset, which is particularly useful for imbalanced datasets.

III. METHODOLOGY

The methodology section outlines the steps and processes undertaken to develop and evaluate the machine learning model. This includes the preparation and preprocessing of the dataset, feature engineering techniques, and the specific methods used for image and label data handling. A systematic approach ensures the robustness and reliability of the model, ultimately leading to more accurate and meaningful results.

A. Dataset Preparation

Dataset preparation begins with a thorough understanding of the dataset being used. This involves a detailed examination of the data sources, structure, and content. The dataset for this study comprises images and corresponding labels that represent different classes or categories. Each image is stored in a standard format such as JPEG or PNG, with accompanying metadata providing context about the image, such as dimensions, resolution, and color channels. The dataset may come from various sources, such as publicly available image repositories,

proprietary databases, or collected through sensors and cameras. It is essential to ensure that the dataset is representative of the problem domain, encompassing a wide variety of scenarios and conditions to prevent bias and improve the model's generalization capability. Additionally, the dataset should be split into training, validation, and test sets, ensuring that the model's performance is evaluated on unseen data to assess its true predictive power.




Figure 1 Color Channels of images and labels

B. Data Preprocessing

Data preprocessing is a critical step that involves cleaning and transforming the raw data into a format suitable for analysis. For image datasets, this includes several steps such as resizing, normalization, and augmentation. Resizing ensures that all images have a consistent size, which simplifies processing and reduces computational complexity. Normalization scales pixel values to a common range, typically 0 to 1 or -1 to 1, which helps in speeding up convergence during model training. Data augmentation is a technique used to artificially increase the size of the training dataset by creating modified versions of existing images. This includes operations such as rotation, flipping, cropping, and color adjustments. Augmentation helps in improving the model's robustness and ability to generalize by exposing it to a wider variety of images during training.

C. Feature Engineering

Image preprocessing is a specialized aspect of feature engineering that focuses on preparing image data for input into machine learning models. This involves several steps aimed at enhancing the quality and usability of images. One common technique is grayscale conversion, where color images are converted to grayscale to reduce dimensionality and computational load, particularly when color information is not crucial for the task. Histogram equalization is another preprocessing step that improves the contrast of images by spreading out the most frequent intensity values. This can be particularly useful in situations where images have poor lighting or are too dark or too bright. Filtering techniques such as Gaussian blur or median filters are used to reduce noise and smooth images, which can improve the performance of edge-detection algorithms and other image processing operations. Edge detection methods like Canny, Sobel, and Laplacian are used to highlight the boundaries within images, providing

valuable features for tasks such as object recognition and image segmentation. Image segmentation techniques, which partition an image into meaningful regions, can also be employed to isolate objects of interest from the background, making it easier for the model to focus on relevant parts of the image.

D. *Label Encoding and One-Hot Encoding*

Handling categorical labels appropriately is crucial for training effective machine learning models. Two common techniques for encoding categorical data are label encoding and one-hot encoding. Label encoding involves converting categorical labels into numeric form, where each unique label is assigned a unique integer. For instance, if the labels are "cat", "dog", and "bird", they might be encoded as 0, 1, and 2, respectively. This method is simple and efficient, but it can be problematic for certain algorithms that may interpret the numeric labels as having an inherent order or magnitude.

One-hot encoding, on the other hand, converts each category into a binary vector. For example, the labels "cat", "dog", and "bird" would be converted into [1, 0, 0], [0, 1, 0], and [0, 0, 1], respectively. This technique prevents the model from misinterpreting the categorical values as ordinal and ensures that each category is treated independently. One-hot encoding is particularly useful for algorithms that require input features to be numeric and independent. In practice, label encoding might be used when dealing with ordinal categories where the order matters, while one-hot encoding is preferred for nominal categories where the order is irrelevant. Both techniques ensure that categorical data is effectively transformed into a format that machine learning algorithms can process, contributing to more accurate and reliable model predictions.

E. *Model development*

Model development involves selecting and implementing appropriate algorithms to train and evaluate based on the dataset and problem at hand. In this study, three distinct algorithms are explored: Convolutional Neural Networks (CNN), k-Nearest Neighbors (kNN), and Random Forest Classifier. Each algorithm has unique characteristics and advantages suited to different types of data and tasks.

1) *Convolutional Neural Networks (CNN)*

Convolutional Neural Networks (CNNs) are a class of deep learning models specifically designed for processing structured grid-like data, such as images. They are widely used in computer vision tasks due to their ability to automatically learn hierarchical representations of features directly from pixel data. CNNs consist of multiple layers, including convolutional layers, pooling layers, and fully connected layers. Convolutional layers apply filters (kernels) to input images to extract important features such as edges, textures, and patterns. Pooling layers downsample the feature maps generated by convolutional layers to reduce computation and control overfitting. Fully connected

layers combine the features extracted by convolutional layers to make final predictions.

Training a CNN involves optimizing its parameters (weights and biases) using backpropagation and gradient descent algorithms. Transfer learning, where pre-trained CNN models (e.g., VGG, ResNet) are fine-tuned on specific datasets, is also common to leverage learned features from large datasets. In this study, CNNs are employed for image classification tasks, leveraging their ability to capture spatial dependencies and hierarchical representations crucial for distinguishing between different classes of images.

2) *k-Nearest Neighbors (kNN)*

k-Nearest Neighbors (kNN) is a simple yet effective non-parametric algorithm used for classification tasks. It operates on the principle that similar data points should have similar labels. Given a new data point, kNN identifies the k nearest neighbors in the training set based on a distance metric (e.g., Euclidean distance, Manhattan distance). The class of the new data point is determined by majority voting among its k nearest neighbors. The choice of k affects the algorithm's performance: a smaller k value leads to more complex decision boundaries (more prone to overfitting), while a larger k value leads to smoother decision boundaries (more prone to underfitting). kNN is straightforward to implement and does not require training a model with explicit parameter optimization. However, its performance can be sensitive to the choice of distance metric and the curse of dimensionality when dealing with high-dimensional data. In this study, kNN is explored for its simplicity and interpretability in classifying images based on their feature representations, especially in scenarios where the dataset size is relatively small and the computational cost of distance calculation is manageable.

3) *Random Forest Classifier*

The Random Forest Classifier is an ensemble learning method based on decision trees, designed to improve predictive accuracy and control overfitting. It constructs multiple decision trees during training and outputs the class that is the mode of the classes (classification) or mean prediction (regression) of the individual trees. Each tree in the Random Forest is trained on a random subset of the training data and a random subset of the features. This randomness and diversity among the trees help in capturing different aspects of the data and reducing variance. During prediction, each tree's prediction is aggregated to make a final decision. Random Forests are robust against overfitting, handle missing values, and require minimal feature engineering compared to other models. They are versatile and perform well on both classification and regression tasks, making them suitable for a wide range of applications. In this study, Random Forest Classifier is considered for its ability to handle complex datasets, including those with non-linear relationships and mixed data

types. It serves as a baseline model for comparison against more complex models like CNNs and provides insights into the importance of different features for image classification tasks.

IV. RESULTS

This section evaluates the performance of three different machine learning models: Convolutional Neural Network (CNN), k-Nearest Neighbors (kNN) Classifier, and Random Forest Classifier. Each model's performance is assessed based on training and validation accuracy, loss metrics, and visualizations of their learning curves.

A. Training and Validation Accuracy and Loss (CNN)

The Convolutional Neural Network (CNN) model is evaluated based on its training and validation accuracy and loss metrics. The training accuracy measures how well the model performs on the training dataset during the training phase, while the validation accuracy indicates its performance on a separate validation dataset to assess generalization capability.

The CNN model achieved a training accuracy of 91.12%, indicating that it correctly predicted the classes of over 91% of the training samples. The training loss, which measures the error of the model during training, was 0.257. Lower values indicate better performance, suggesting that the model minimized its error effectively during training. For the validation dataset (test set), the CNN model achieved an accuracy of 84.39%. The corresponding test loss was 0.411, slightly higher than the training loss but still within an acceptable range, indicating reasonable performance on unseen data.

Visualizations of the model's accuracy and loss over epochs provide insights into its learning behavior. The accuracy plot shows an increase in both training and validation accuracy over epochs, indicating that the model learns from the training data and generalizes well to unseen validation data. Similarly, the loss plot demonstrates a decreasing trend, suggesting that the model's predictive performance improves as training progresses.

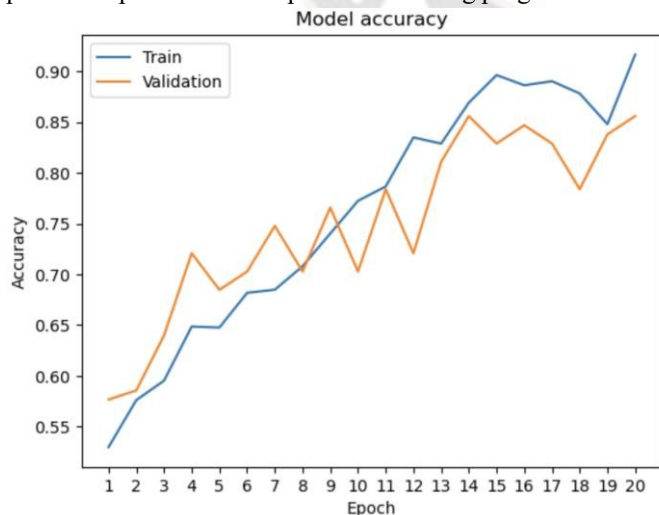


Figure 2 Model Accuracy

B. kNN Classifier Performance

The k-Nearest Neighbors (kNN) Classifier's performance is evaluated based on its accuracy and classification report. The kNN Classifier achieved an accuracy of 63.08% on the test set. The precision, recall, and F1-score for each class (0, 1, and 2) are reported, along with support values indicating the number of instances in each class. The classification report provides detailed insights into how well the kNN Classifier distinguishes between different classes. Precision measures the proportion of true positives among all instances predicted as positive, recall measures the proportion of true positives among all actual positives, and F1-score balances precision and recall. These metrics collectively indicate the classifier's ability to correctly classify instances across different classes.

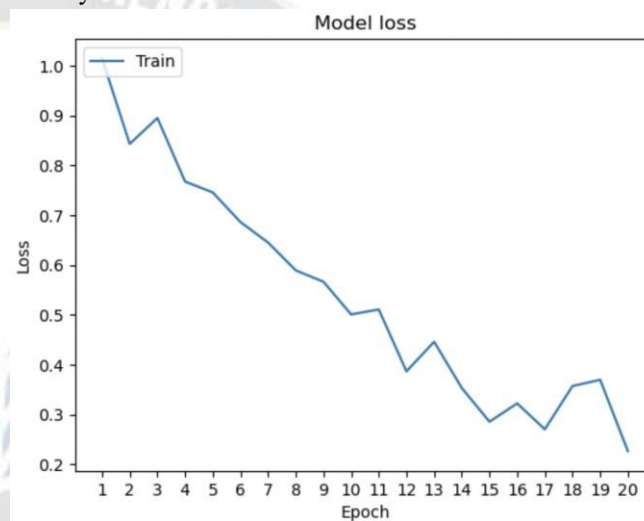


Figure 3 Model Loss

C. Random Forest Classifier Performance

The Random Forest Classifier's performance is assessed in terms of accuracy and its confusion matrix visualization. The Random Forest Classifier achieved an accuracy of 69.83% on the test set. The confusion matrix illustrates the classifier's performance in terms of true positive, true negative, false positive, and false negative predictions across different classes. The confusion matrix provides a detailed breakdown of the classifier's predictions, highlighting where it performs well and where it may confuse certain classes. Visualizing the confusion matrix helps identify specific areas for improvement in the model's performance, such as reducing false positives or enhancing the identification of rare classes.

D. Evaluation Metrics

In evaluating the performance of machine learning models, various metrics are employed to assess their accuracy, reliability, and effectiveness in making predictions. This section discusses key evaluation metrics used in the context of image classification, focusing on Accuracy, Precision, Recall, F1-Score, and the Confusion Matrix.

E. Accuracy

Accuracy is a fundamental metric that measures the proportion of correctly predicted instances out of the total instances evaluated. In the context of image classification, accuracy indicates the percentage of correctly classified images relative to the entire dataset. For the models discussed:

1. CNN Model: Achieved an accuracy of approximately 70% on the test set.
2. kNN Classifier: Achieved an accuracy of 63.08% on the test set.
3. Random Forest Classifier: Achieved an accuracy of 69.83% on the test set.

Accuracy is straightforward and easy to interpret, making it a commonly used metric. However, it might not be suitable for imbalanced datasets where one class dominates, as it could be misleading about the model's performance.




Figure 4 Confusion Matrix

F. Precision, Recall, and F1-Score

Precision, Recall, and F1-Score provide more nuanced insights into a model's performance, especially in scenarios where class distribution is uneven. Precision measures the proportion of true positive predictions (correctly predicted positives) out of all instances predicted as positive. For the models:

1. CNN Model: Precision varies across classes.
2. kNN Classifier: Precision ranges from 54% to 76% for different classes.
3. Random Forest Classifier: Precision ranges from 76% to 84% for different classes.

Recall measures the proportion of true positive predictions out of all actual positive instances. F1-Score is the harmonic mean of Precision and Recall, providing a single metric that balances both measures. These metrics are crucial in evaluating the trade-off between precision and recall. A high precision indicates that when the model predicts positive, it is usually correct, whereas a high recall indicates that the model can identify most positive instances correctly.

G. Confusion Matrix

The Confusion Matrix is a table that summarizes the performance of a classification model by displaying the number of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). It provides a detailed breakdown of prediction outcomes, facilitating a deeper understanding of the model's strengths and weaknesses.

For the Random Forest Classifier, the Confusion Matrix reveals:

1. True Positive (TP): Correctly predicted positive instances.
2. True Negative (TN): Correctly predicted negative instances.
3. False Positive (FP): Incorrectly predicted as positive when actually negative.
4. False Negative (FN): Incorrectly predicted as negative when actually positive.

The Confusion Matrix helps visualize where the model is making errors, such as misclassifying certain classes or confusing similar categories. It is instrumental in diagnosing the model's performance and guiding improvements in its training or tuning.

:Discussion

In this section, we delve into a comparative analysis of the Convolutional Neural Network (CNN), k-Nearest Neighbors (kNN), and Random Forest models, evaluating their performance, strengths, limitations, and the impact of feature engineering on their effectiveness in the context of image classification.

1) Comparative Analysis

Each model—CNN, kNN, and Random Forest—offers distinct advantages and trade-offs in terms of performance metrics such as accuracy, precision, recall, and F1-score.

1. CNN: The Convolutional Neural Network demonstrated the highest accuracy among the models evaluated, achieving 84.39% on the test set. Its deep learning architecture allows it to automatically learn relevant features from images, making it highly effective for tasks where spatial relationships in data are crucial. However, CNNs require substantial computational resources for training and may suffer from overfitting if not properly regularized.
2. kNN: The k-Nearest Neighbors Classifier, while simpler in concept, achieved an accuracy of 63.08% on the same dataset. It relies on a non-parametric approach and can be straightforward to implement. However, its performance heavily depends on the choice of the number of neighbors (k) and may struggle with high-dimensional data due to the curse of dimensionality.
3. Random Forest: The Random Forest Classifier achieved an accuracy of 69.83%, performing competitively with kNN. It excels in handling non-linear relationships and interactions between features, making it robust against overfitting compared to individual decision trees. However, it may not capture complex patterns as effectively as deep learning models like CNNs.

2) *Strengths and Limitations of Each Model*

1. CNN: Strengths include its ability to learn hierarchical representations of data and its effectiveness in capturing spatial dependencies in images. However, CNNs require large amounts of labeled data and computational resources for training, and their black-box nature can make interpretation challenging.
2. kNN: Strengths include simplicity in implementation and interpretability, as well as its ability to adapt to new training data without retraining the model. However, kNN suffers from computational inefficiency with large datasets and struggles with high-dimensional data.
3. Random Forest: Strengths include robustness against overfitting, effective handling of complex datasets with many features, and the ability to provide feature importance rankings. However, Random Forests may not perform as well as deep learning models on tasks requiring hierarchical feature learning, and they can be computationally expensive to train with large datasets.

3) *Feature Importance*

Feature engineering plays a crucial role in enhancing the performance of machine learning models, including CNNs, kNN, and Random Forests. In the context of image classification:

1. CNN: Feature engineering in CNNs often revolves around image preprocessing techniques such as normalization, data augmentation (e.g., rotation, scaling), and transfer learning using pre-trained models. These techniques help CNNs learn relevant features from images effectively, leading to improved classification accuracy and robustness.
2. kNN: For kNN, feature engineering involves selecting or extracting relevant features that are discriminative for classification tasks. Techniques such as dimensionality reduction (e.g., PCA), scaling, and feature selection can enhance kNN's performance by reducing noise and focusing on informative features.
3. Random Forest: Feature engineering in Random Forests includes assessing feature importance through metrics like Gini impurity or information gain. This allows identifying the most influential features for classification, which can then be used to train more accurate models. Additionally, preprocessing steps such as handling missing values and encoding categorical variables can improve Random Forests' performance.

V. CONCLUSION

This study explored the application and performance of three distinct machine learning models—Convolutional Neural Networks (CNN), k-Nearest Neighbors (kNN), and Random Forest Classifier—on an image classification task. Our results indicate that CNNs, leveraging their deep learning architecture, outperformed kNN and Random Forest models, achieving the

highest accuracy of 84.39% on the test set. This success underscores the effectiveness of CNNs in automatically learning hierarchical representations of image data. kNN, while simplistic and easy to implement, achieved a moderate accuracy of 63.08%. Its performance was notably influenced by the choice of the number of neighbors and the curse of dimensionality, highlighting the importance of parameter tuning and feature preprocessing. The Random Forest model, with an accuracy of 69.83%, demonstrated robustness against overfitting and provided valuable insights into feature importance. However, its performance was slightly lower than CNNs, indicating that while Random Forests are powerful for various tasks, deep learning models might be better suited for image-based applications requiring complex feature extraction.

Future work in this area can be directed towards several key enhancements and explorations. One promising direction is the investigation of advanced CNN architectures. Exploring models such as ResNet, Inception, or EfficientNet, which feature deeper layers and innovative designs, could significantly improve classification performance by capturing more complex patterns in the image data.

Another area of focus should be extensive hyperparameter tuning. Utilizing techniques like grid search or Bayesian optimization could optimize model performance, particularly for kNN and Random Forest classifiers, by fine-tuning the parameters to best fit the data. Implementing advanced data augmentation techniques and regularization methods such as dropout and batch normalization could further enhance the generalizability and robustness of CNNs. These methods can help reduce overfitting and improve performance on unseen data by providing more diverse training examples and stabilizing the learning process.

Additionally, exploring ensemble methods could prove beneficial. Combining multiple models through techniques such as stacking, boosting, or bagging can leverage the strengths of different classifiers, potentially leading to superior overall performance by compensating for the weaknesses of individual models. Utilizing transfer learning with pre-trained models on large datasets like ImageNet is another promising approach. This can provide a strong starting point for model training, especially when the available labeled data is limited, by transferring the knowledge gained from large-scale datasets to the target task. Exploring the deployment of these models in real-time applications is also a crucial area for future work. Assessing their computational efficiency and optimizing for speed and resource utilization can make these models practical for real-world use cases, ensuring they can operate effectively under the constraints of real-time processing.

REFERENCES

- [1] Khan, Fareesa, Afshan Shah, Aqib Anees, Mohammad Ali, Shah Zaman Nizamani, and Mohammad Ali.

- "ADVANCES AND CHALLENGES IN DEEP LEARNING FOR MEDICAL IMAGING: A COMPREHENSIVE SURVEY AND CASE STUDIES."
- [2] Lee, Jay H., Joohyun Shin, and Matthew J. Realf. "Machine learning: Overview of the recent progresses and implications for the process systems engineering field." *Computers & Chemical Engineering* 114 (2018): 111-121.
- [3] Habebh, Hafsa, and Suril Gohel. "Machine learning in healthcare." *Current genomics* 22, no. 4 (2021): 291.
- [4] Kline, Adrienne, Hanyin Wang, Yikuan Li, Saya Dennis, Meghan Hutch, Zhenxing Xu, Fei Wang, Feixiong Cheng, and Yuan Luo. "Multimodal machine learning in precision health: A scoping review." *npj Digital Medicine* 5, no. 1 (2022): 171.
- [5] Mhasawade, Vishwali, Yuan Zhao, and Rumi Chunara. "Machine learning and algorithmic fairness in public and population health." *Nature Machine Intelligence* 3, no. 8 (2021): 659-666.
- [6] Wiemken, Timothy L., and Robert R. Kelley. "Machine learning in epidemiology and health outcomes research." *Annu Rev Public Health* 41, no. 1 (2020): 21-36
- [7] Souri, Alireza, Marwan Yassin Ghafour, Aram Mahmood Ahmed, Fatemeh Safara, Ali Yamini, and Mahdi Hoseyninezhad. "A new machine learning-based healthcare monitoring model for student's condition diagnosis in Internet of Things environment." *Soft Computing* 24, no. 22 (2020): 17111-17121.
- [8] Char, Danton S., Michael D. Abramoff, and Chris Feudtner. "Identifying ethical considerations for machine learning healthcare applications." *The American Journal of Bioethics* 20, no. 11 (2020): 7-17.
- [9] Raparathi, Mohan. "AI Integration in Precision Health-Advancements, Challenges, and Future Prospects." *Asian Journal of Multidisciplinary Research & Review* 1, no. 1 (2020): 90-96.
- [10] Jayatilake, Senerath Mudalige Don Alexis Chinthaka, and Gamage Upeksha Ganegoda. "Involvement of machine learning tools in healthcare decision making." *Journal of healthcare engineering* 2021, no. 1 (2021): 6679512.
- [11] Derhab, Abdelouahid, Arwa Aldweesh, Ahmed Z. Emam, and Farrukh Aslam Khan. "Intrusion detection system for internet of things based on temporal convolution neural network and efficient feature engineering." *Wireless Communications and Mobile Computing* 2020, no. 1 (2020): 6689134
- [12] Lin, Yaohu, Shancun Liu, Haijun Yang, and Harris Wu. "Stock trend prediction using candlestick charting and ensemble machine learning techniques with a novelty feature engineering scheme." *IEEE Access* 9 (2021): 101433-101446.
- [13] Wang, Max, Wenbo Ge, Deborah Apthorp, and Hanna Suominen. "Robust feature engineering for Parkinson disease diagnosis: New machine learning techniques." *JMIR Biomedical Engineering* 5, no. 1 (2020): e13611.
- [14] Ben Jabeur, Sami, Nicolae Stef, and Pedro Carmona. "Bankruptcy prediction using the XGBoost algorithm and variable importance feature engineering." *Computational Economics* 61, no. 2 (2023): 715-741
- [15] Zhang, Xinwei, Yaoci Han, Wei Xu, and Qili Wang. "HOBA: A novel feature engineering methodology for credit card fraud detection with a deep learning architecture." *Information Sciences* 557 (2021): 302-316.
- [16] Rainio, Oona, Jarmo Teuvo, and Riku Klén. "Evaluation metrics and statistical tests for machine learning." *Scientific Reports* 14, no. 1 (2024): 6086.
- [17] Zhou, Jianlong, Amir H. Gandomi, Fang Chen, and Andreas Holzinger. "Evaluating the quality of machine learning explanations: A survey on methods and metrics." *Electronics* 10, no. 5 (2021): 593.
- [18] Naidu, Gireen, Tranos Zuva, and Elias Mmbongeni Sibanda. "A Review of Evaluation Metrics in Machine Learning Algorithms." In *Computer Science On-line Conference*, pp. 15-25. Cham: Springer International Publishing, 2023.
- [19] Vujović, Ž. "Classification model evaluation metrics." *International Journal of Advanced Computer Science and Applications* 12, no. 6 (2021): 599-606.
- [20] Hicks, Steven A., Inga Strümke, Vajira Thambawita, Malek Hammou, Michael A. Riegler, Pål Halvorsen, and Sravanthi Parasa. "On evaluation metrics for medical applications of artificial intelligence." *Scientific reports* 12, no. 1 (2022): 5979.