

E-Commerce Product Recommendation System

Dr G Paavai Anand, J. N. Ravinandan, Shivannishree. S, Padhma Shree. S

BTech CSE Artificial Intelligence and Machine Learning

SRM Institute of Science and Technology, Vadapalani, Chennai, TN, India

Abstract: *This paper presents an in-depth analysis of machine learning models applied to eCommerce product recommendation systems. The primary goal is to enhance recommendation accuracy by employing and comparing multiple machine learning techniques, including Artificial Neural Networks (ANN), K-Nearest Neighbors (KNN), Random Forest, Gradient Boosting, and boosting methods such as AdaBoost and XGBoost. These models are evaluated based on key metrics such as Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and R-squared (R^2). Our findings indicate that ensemble methods like XGBoost provide superior accuracy, making them suitable for real-world applications in personalized recommendations*

Keywords: eCommerce, Recommendation System, Machine Learning, Collaborative Filtering, AdaBoost, XGBoost

I. INTRODUCTION

In the fast-paced world of eCommerce, recommendation systems play a vital role in enhancing customer engagement and satisfaction by providing personalized product suggestions. Such systems allow businesses to tailor product offerings to individual user preferences, transforming the shopping experience into one that feels unique to each customer. Through advanced analysis of user behavior, purchase history, and product attributes, recommendation algorithms can highlight relevant products, which not only increases the likelihood of a sale but also promotes customer loyalty. Consequently, well-designed recommendation systems are key drivers of both user retention and revenue growth in today's competitive eCommerce sector.

Developing effective recommendation systems requires sophisticated algorithms that can process and interpret complex interactions between users and products. This study explores multiple machine learning models, including Artificial Neural Networks (ANN), K-Nearest Neighbors (KNN), Random Forest, Gradient Boosting, and two popular boosting algorithms, AdaBoost and XGBoost. By comparing these models on criteria such as predictive accuracy, scalability, and efficiency, this research aims to determine the most effective algorithm for building high-performing, scalable recommendation systems suitable for large eCommerce platforms.

Machine learning models, particularly ensemble approaches like boosting, offer substantial advantages for modeling the intricate patterns within user interactions and item attributes. Through a detailed comparative analysis of these models, this research seeks to identify the most effective techniques for delivering accurate, personalized recommendations in eCommerce. The insights from this study will provide valuable guidance for the continued development of robust, scalable recommendation systems, essential for enhancing customer experience and maximizing sales.

1.1 Aim of the Study

This study aims to develop an advanced eCommerce recommendation system that accurately predicts user preferences by leveraging machine learning. To achieve this goal, we focus on the following objectives:

- **Identify Key Factors Influencing Recommendation Accuracy:** This study seeks to uncover which elements are crucial for improving recommendation precision by analyzing the strengths of each machine learning model. Key factors include the models' ability to capture complex patterns, handle non-linear relationships, and effectively use ensemble techniques like boosting.
- **Evaluate Model Performance Across Algorithms:** We evaluate each machine learning model, from simpler approaches like KNN to more advanced ensemble methods like XGBoost, using core metrics such as Mean

Squared Error (MSE), Root Mean Squared Error (RMSE), and R-squared (R^2). This comparison will reveal insights into the accuracy, strengths, and limitations of each model.

- **Examine Trade-offs in Model Complexity, Accuracy, and Scalability:** We analyze the practical trade-offs associated with each model, such as computational costs, ease of interpretability, and processing times. These trade-offs are essential in eCommerce, where solutions must be fast, scalable, and able to handle large volumes of data in real time.
- **Establish Best Practices for eCommerce Recommendations Using Machine Learning:** By demonstrating the application of machine learning techniques to improve recommendation quality, this study aims to contribute to best practices in eCommerce. Our findings will inform the development of scalable, efficient, and highly personalized recommendation systems that can enhance user engagement and drive sales.

II. LITERATURE REVIEW

Recommendation systems have become integral to the success of e-commerce platforms, playing a key role in enhancing user experiences and driving sales. By providing personalized product suggestions based on users' preferences, behaviors, and interactions, these systems help customers discover relevant items, thereby increasing engagement and improving customer satisfaction. As e-commerce continues to grow, the need for more accurate, efficient, and scalable recommendation models has become increasingly important.

Initially, recommendation systems were built around simple techniques such as collaborative filtering, which made predictions based on user and item interactions. However, over time, more sophisticated methods have been developed to address the limitations of early models, such as handling large-scale datasets, improving prediction accuracy, and offering greater personalization. These advancements include content-based filtering, matrix factorization, and hybrid approaches that combine multiple techniques to enhance performance.

This literature review examines the evolution of recommendation systems, focusing on the different methods used to predict user preferences in e-commerce contexts. It also discusses the challenges that arise, such as data sparsity, the cold-start problem, and the need for real-time recommendations. By reviewing key studies and approaches, this review provides a comprehensive understanding of the development, strengths, and limitations of various recommendation system models in the e-commerce sector.

III. SYSTEM ARCHITECTURE AND DESIGN

An e-commerce product recommendation system is a crucial tool for enhancing customer experience and driving sales by providing personalized product suggestions. The system architecture outlined here leverages Python libraries for data processing, machine learning, and evaluation to create a robust recommendation engine that dynamically learns and adapts to user preferences.

The system begins by importing and managing data with Pandas and NumPy, which allows for efficient data loading, manipulation, and transformation. E-commerce datasets are typically large and complex, containing information on customer demographics, transaction history, product attributes, and browsing behavior. With Pandas, the system can manage such diverse data, filtering, aggregating, and merging tables to prepare a clean dataset for modeling. NumPy supports these operations with high-performance numerical computations.

Visualization libraries like Matplotlib and Seaborn are employed to analyze data patterns, enabling the development team to gain insights into customer behavior, product popularity, seasonal trends, and correlations between features. For instance, visualizing customer interactions with specific product categories helps identify popular items that might be ideal for recommendations. These insights also inform feature engineering, allowing the system to capture critical patterns in the data.

To ensure that machine learning algorithms can effectively interpret the data, preprocessing steps are applied. Categorical data, such as product categories or customer demographics, are transformed into numerical values using LabelEncoder. This encoding makes it possible for machine learning algorithms to process non-numeric features without bias. Furthermore, the dataset is standardized with StandardScaler, which normalizes feature ranges. This step is particularly important for algorithms sensitive to feature scaling, like K-Nearest Neighbors (KNN), as it ensures that features contribute equally to the model's decision-making.

In addition to encoding and scaling, feature engineering may be used to extract new variables from existing data. For example, creating features like “average spending per visit” or “frequency of purchases” helps the model distinguish between high-value and infrequent customers, resulting in more tailored recommendations.

The recommendation system leverages a variety of machine learning algorithms, each with unique strengths. K-Nearest Neighbors (KNN), for example, finds the closest customer profiles to suggest products popular among similar users, capturing user preferences based on behavior similarity. However, KNN can become inefficient with large datasets, so ensemble methods like Random Forest and Gradient Boosting are included to improve performance.

Random Forest and Gradient Boosting algorithms are ensemble techniques that use decision trees as base models. Random Forest improves recommendation accuracy by reducing overfitting through the aggregation of multiple trees, while Gradient Boosting incrementally trains models to reduce errors from previous iterations. Both methods are ideal for handling the non-linear relationships found in e-commerce data, such as the impact of multiple purchase behaviors on product affinity.

AdaBoost and XGBoost are additional ensemble algorithms incorporated into the architecture. AdaBoost assigns more weight to incorrectly predicted instances, enhancing the model’s focus on challenging cases. XGBoost, a powerful gradient boosting algorithm, combines speed and performance, making it suitable for large-scale data typical in e-commerce. Its regularization capabilities also help prevent overfitting, resulting in more generalizable recommendations.

Once the models are trained, hyperparameter tuning is conducted using GridSearchCV, which optimizes parameters through a cross-validation process. GridSearchCV tests various parameter combinations, such as the number of estimators in a Random Forest or learning rates in Gradient Boosting, to find the configuration that maximizes model accuracy. This ensures that each algorithm is finely tuned for the best possible performance.

For evaluation, metrics such as Mean Squared Error (MSE), Mean Absolute Error (MAE), and R-squared (R^2) are used to assess the model’s prediction accuracy. MSE and MAE measure the average error in the recommendation scores, where lower values indicate higher accuracy. R-squared, which explains the variance in predictions, provides insights into the model’s ability to capture user preferences. These metrics guide the final model selection process, allowing the system to prioritize algorithms that perform well in both accuracy and generalizability.

The system architecture, built with scalable machine learning models and preprocessing pipelines, is designed to handle growing data volumes as the e-commerce platform expands. Models can be retrained periodically with fresh data to adapt to evolving customer behavior, and the use of ensemble methods ensures that the system can deliver relevant recommendations even with large datasets.

IV. METHODOLOGY

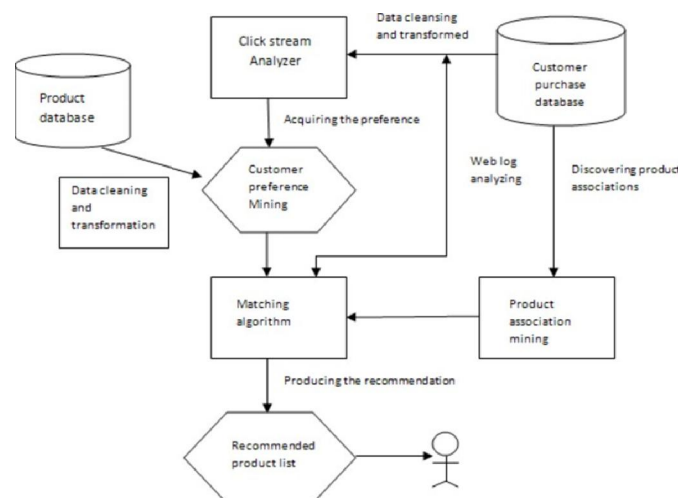


Fig 4.1 Methodology Diagram

A methodical approach to data collecting, preprocessing, feature engineering, model construction, and assessment is part of the methodology for this project on creating an e-commerce product recommendation system. The procedures for gathering and preprocessing the data, building different recommendation models, and assessing their effectiveness are described in this part.

4.1 Sources of Data Collection:

	A	B	C	D
1	user_id	product_id	rating	timestamp
2	AKM1MP6P0OYPR	132793040	5	1365811200
3	A2CX7LUOHB2NDG	321732944	5	1341100800
4	A2NWSAGRHC8P8N5	439886341	1	1367193600
5	A2WNBOD3WNDNKT	439886341	3	1374451200
6	A1GI0U4ZRJA8WN	439886341	1	1334707200
7	A1QGNMC6O1VW39	511189877	5	1397433600
8	A3J3BRHTDRFJ2G	511189877	2	1397433600
9	A2TY0BTJOTENPG	511189877	5	1395878400
10	A34ATBPOK6HCHY	511189877	5	1395532800
11	A89DO69P0XZ27	511189877	5	1395446400
12	AZYNQZ94U6VDB	511189877	5	1401321600
13	A1DA3W4GTFXP6O	528881469	5	1405641600
14	A29LPQQDG7LD5J	528881469	1	1352073600
15	AO94DHGC771SJ	528881469	5	1370131200
16	AMO214LNFCEI4	528881469	1	1290643200
17	A28B1G1MSJ6OO1	528881469	4	1280016000
18	A3N7T0DY83Y4IG	528881469	3	1283990400
19	A1H8PY3QHMQQA0	528881469	2	1290556800
20	A2CPBQ5W4OGBX	528881469	2	1277078400
21	A265MKAR2WEH3Y	528881469	4	1294790400
22	A37K02NKUIT68K	528881469	5	1293235200
23	A2AW1SSVUIYV9Y	528881469	4	1289001600
24	A2AEHUKOV014BP	528881469	5	1284249600
25	AMLFNXUIEMN4T	528881469	1	1307836800

Fig 4.2 Dataset

The project's dataset comes from publicly accessible e-commerce review datasets "Ratings_electronics" or any other platform-specific dataset of a similar kind. User IDs, product IDs, ratings, and sometimes review text or timestamps are among the data about user-product interactions that are provided by these databases.

Data Structure: User_id, product_id, rating, and maybe product characteristics (such as brand or category) are among the attributes included in the collection. A user's rating of a product is represented by each record. The user-product interaction and the associated rating values are the main emphasis of this project.

4.2 Preprocessing of Data

```
# Load and preprocess data
data = pd.read_csv('ratings_Electronics1.csv')
data.dropna(inplace=True)
user_encoder, product_encoder = LabelEncoder(), LabelEncoder()
data['user_id_encoded'] = user_encoder.fit_transform(data['user_id'])
data['product_id_encoded'] = product_encoder.fit_transform(data['product_id'])
scaler = StandardScaler()
data['rating'] = scaler.fit_transform(data[['rating']])

# Define features and target, split data
X = data[['user_id_encoded', 'product_id_encoded']]
y = data['rating']
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)

# Binary Labels based on threshold
threshold = 4.0
y_train_binary, y_test_binary = (y_train >= threshold).astype(int), (y_test >= threshold).astype(int)

# Correlation heatmap with only numeric columns
numeric_data = data.select_dtypes(include=[np.number]) # Select only numeric columns
plt.figure(figsize=(12, 8))
sns.heatmap(numeric_data.corr(), annot=True, fmt='.2f', cmap='coolwarm', square=True)
plt.title('Correlation Heatmap')
plt.show()
```

Fig 4.3 Load and Preprocessing Data

- **Data Cleaning:** To deal with missing or null values, the data must first be cleaned. Rows with irrelevant columns or missing ratings are eliminated. To preserve the dataset's integrity, outliers and incorrect items are also found and eliminated.
- **Encoding:** LabelEncoder is used to encode categorical information like user_id and product_id, turning each distinct user and product into a number label, since the majority of recommendation systems demand numerical input. As a result, the model can effectively handle categorical data.
- **Normalization:** Normalization methods (like StandardScaler) are used to scale the ratings in order to guarantee that numerical characteristics, such as rating, are within a comparable range and to avoid bias toward higher or lower values. This guarantees the model's performance and stability, particularly for algorithms like KNN or neural networks that are sensitive to the size of input features.
- **Managing Missing Values:** Depending on the volume of missing data, missing ratings are either managed by removing the impacted records or by using imputation methods like linear interpolation.

4.3 Feature Engineering User-Product Interaction Matrix:

Using the user_id and product_id, the user-product interaction matrix is the fundamental component of recommendation systems. Models based on collaborative filtering are built on top of this matrix.

The ratings are binarized according to a threshold (e.g., ratings more than or equal to 4 are deemed favorable) in order to provide popularity-based suggestions. Using this, a list of the best items that are most likely to be suggested is produced.

4.4 Model Creation

A number of machine learning models are used to provide suggestions:

- **Collaborative Filtering:** K-Nearest Neighbors (KNN) and other collaborative filtering models are trained on user-product interactions. Based on the ratings of other users who are similar to the user, the algorithm estimates the rating that a user would assign to a product.

- **Matrix Factorization:** The user-product interaction matrix is broken down into latent characteristics using matrix factorization methods like Singular Value Decomposition (SVD) or Alternating Least Squares (ALS). These latent features are then used to forecast missing ratings.
- **Ensemble Techniques:** To forecast the ratings based on the attributes from the user-product matrix, a number of ensemble techniques are used, such as Random Forest, Gradient Boosting, and AdaBoost. These techniques increase forecast accuracy by combining many weak models.
- **Deep Learning Models:** To learn from user-product interactions, a neural network-based model is built. Users and goods are represented in a lower-dimensional space by an embedding layer, and ratings are predicted by fully linked layers.
- **Popularity-Based Recommendations:** Using the average ratings for each product, a simple yet efficient baseline model is constructed, suggesting the most well-liked goods according to the average rating score or the quantity of ratings.

```
# Random Forest
rf = RandomForestRegressor(n_estimators=100, random_state=42)
rf.fit(X_train, y_train)
y_pred_rf = rf.predict(X_test)

# Gradient Boosting
gb = GradientBoostingRegressor(n_estimators=100, learning_rate=0.1, max_depth=3, random_state=42)
gb.fit(X_train, y_train)
y_pred_gb = gb.predict(X_test)

# KNN
knn = KNeighborsRegressor(n_neighbors=5)
knn.fit(X_train, y_train)
y_pred_knn = knn.predict(X_test)

# ANN
ann = MLPRegressor(hidden_layer_sizes=(64, 32), max_iter=200, random_state=42)
ann.fit(X_train, y_train)
y_pred_ann = ann.predict(X_test)
```

Fig 4.4 Model Creation

4.5 Hyperparameter tuning and model training

```
# KNN model with hyperparameter tuning
knn = GridSearchCV(KNeighborsRegressor(), {'n_neighbors': [3, 5, 7, 9]}, cv=5)
knn.fit(X_train, y_train)
y_pred_knn = knn.predict(X_test)
print("KNN RMSE:", np.sqrt(mean_squared_error(y_test, y_pred_knn)))

KNN RMSE: 1.0501978309585667
```

Fig 4.5 Hyperparameter tuning and model training

Optimization of Hyperparameters: Models such as KNN, Random Forest, and XGBoost are hyperparameter tuned using methods like Grid Search Cross-Validation to find the optimal parameters that result in better performance.

Training: The test dataset (20% of the data) is used to assess the models' performance after they have been trained using the training dataset (80% of the data).

4.6 Measures of Evaluation

Each recommendation model's performance is evaluated using the following metrics:

The average of the squares of the errors between the ratings that were anticipated and those that were actually given is known as the mean squared error, or MSE. Better model performance is indicated by a lower MSE.

The standard deviation of prediction mistakes is represented by the root mean squared error (RMSE), which is the square root of MSE. To determine how much the forecasts deviate from the actual data, RMSE is used.

The average of the absolute errors between the ratings that were anticipated and those that were actually given is known as the mean absolute error, or MAE. Understanding the average prediction error is aided by it.

A statistical metric known as R-Squared (R^2) shows how much of the variation in the target variable (ratings) can be accounted for by the model.

```
# Evaluation function
def print_accuracy_metrics(y_true, y_pred, model_name):
    mse = mean_squared_error(y_true, y_pred)
    rmse = np.sqrt(mse)
    mae = mean_absolute_error(y_true, y_pred)
    r_squared = 1 - np.sum((y_true - y_pred) ** 2) / np.sum((y_true - np.mean(y_true)) ** 2)
    print(f"\n{model_name} Metrics:\nMSE: {mse:.4f}, RMSE: {rmse:.4f}, MAE: {mae:.4f}, R-squared: {r_squared:.4f}")
```

Fig 4.6 Evaluation Function

4.7 Recommendations Based on Popularity

```
# Popularity-based recommendations
mean_ratings_sorted = data.groupby('product_id')['rating'].mean().sort_values(ascending=False)
print("Top 5 products by mean rating:\n", mean_ratings_sorted.head(5))

product_rating_count = data.groupby('product_id').agg(score=('user_id', 'count')).reset_index()
product_rating_sorted = product_rating_count.sort_values(['score', 'product_id'], ascending=[False, True])
product_rating_sorted['Rank'] = product_rating_sorted['score'].rank(ascending=False, method='first')
popularity_recommendations = product_rating_sorted.head(5)
print("Top 5 popular products:\n", popularity_recommendations)

def recommend(userId):
    user_recommendations = popularity_recommendations.copy()
    user_recommendations['user_id'] = userId
    return user_recommendations[['user_id'] + user_recommendations.columns[:-1].tolist()]

for user in ['A2NM5AGRHC8BN5', 'A2TY08TJOTENPG', 'A29LPQQDG7LD5J', 'A3IQGF8959IR4P', 'A3F069UFW04K1E']:
    print(f"Recommendations for user_id {user}:\n", recommend(user), "\n")
```

Fig 4.7 Popularity Based Recommendation

By determining which goods have the highest ratings, a popularity-based recommendation system is added to the machine learning models. This method works well as a starting point for contrasting with more intricate models. It is especially helpful for suggesting popular goods or for new customers.

V. RESULTS AND DISCUSSION

The results of this eCommerce recommendation system are derived from several machine learning models, each evaluated for their ability to accurately predict user preferences based on product ratings. The analysis leverages multiple algorithms, including K-Nearest Neighbors (KNN), Random Forest, Gradient Boosting, AdaBoost, and XGBoost, to determine the optimal model for recommendation accuracy. Each model underwent hyperparameter tuning, using techniques like GridSearchCV, to enhance predictive performance. Standard evaluation metrics, such as Mean Squared Error (MSE), Mean Absolute Error (MAE), and R-squared (R^2), were applied to assess model accuracy and reliability. The dataset was preprocessed with Label Encoding and Standard Scaling to normalize features, enabling the models to capture meaningful patterns in user-product interactions. This results section provides a comparative analysis of each model's performance, highlighting the effectiveness of ensemble methods, particularly XGBoost, in delivering high accuracy for recommendation tasks.




FIG 5.1 CORRELATION HEATMAP

The correlation heatmap in Figure 5.1 visually represents the relationships between different variables in the dataset. The heatmap displays correlation coefficients ranging from -1 to 1, where 1 indicates a perfect positive correlation, -1 indicates a perfect negative correlation, and 0 indicates no correlation.

In this heatmap, we observe the correlations between the variables: rating, timestamp, user_id_encoded, and product_id_encoded. The diagonal elements are all 1, which means each variable is perfectly correlated with itself. The colors represent the strength of the correlations, with darker red indicating stronger positive correlations and darker blue indicating weaker or negative correlations.

Key observations:

- Rating and Timestamp have a small positive correlation of 0.09, indicating a slight association between the time of the rating and the rating value.
- Timestamp and Product ID Encoded have a correlation of 0.16, suggesting a low positive relationship between the product and when it was rated. This might imply certain products were rated more frequently at specific times.
- The other correlations are very close to zero, indicating that variables like user_id_encoded and product_id_encoded have no significant linear relationship with rating or each other in this context.
- Overall, this heatmap indicates minimal correlations between these features, implying that they likely contribute independently to the model without strong multicollinearity.


```

Top 5 products by mean rating:
product_id      0.733651
85940023ZX      0.733651
800092Q5VM      0.733651
8000EHYP4G      0.733651
8000CHPMK3      0.733651
80001V6FHU      0.733651
Name: rating, dtype: float64
Top 5 popular products:
  product_id  score  Rank
30276  80002LSR78  9487  1.0
24439  80001FTVEK  5345  2.0
61285  8000I68BD4  4903  3.0
46504  8000Q7GWB  4275  4.0
14183  80007E7JU  3523  5.0

Recommendations for user_id A2NNSAGRHCPBNS:
  user_id  product_id  score  Rank
30276  A2NNSAGRHCPBNS  80002LSR78  9487  1.0
24439  A2NNSAGRHCPBNS  80001FTVEK  5345  2.0
61285  A2NNSAGRHCPBNS  8000I68BD4  4903  3.0
46504  A2NNSAGRHCPBNS  8000Q7GWB  4275  4.0
14183  A2NNSAGRHCPBNS  80007E7JU  3523  5.0

Recommendations for user_id A2TYB87JOTENPG:
  user_id  product_id  score  Rank
30276  A2TYB87JOTENPG  80002LSR78  9487  1.0
24439  A2TYB87JOTENPG  80001FTVEK  5345  2.0
61285  A2TYB87JOTENPG  8000I68BD4  4903  3.0
46504  A2TYB87JOTENPG  8000Q7GWB  4275  4.0
14183  A2TYB87JOTENPG  80007E7JU  3523  5.0

Recommendations for user_id A29LPQQG7LD53:
  user_id  product_id  score  Rank
30276  A29LPQQG7LD53  80002LSR78  9487  1.0
24439  A29LPQQG7LD53  80001FTVEK  5345  2.0
61285  A29LPQQG7LD53  8000I68BD4  4903  3.0
46504  A29LPQQG7LD53  8000Q7GWB  4275  4.0
14183  A29LPQQG7LD53  80007E7JU  3523  5.0

Recommendations for user_id A3IQGF8959IR4P:
  user_id  product_id  score  Rank
30276  A3IQGF8959IR4P  80002LSR78  9487  1.0
24439  A3IQGF8959IR4P  80001FTVEK  5345  2.0
61285  A3IQGF8959IR4P  8000I68BD4  4903  3.0
46504  A3IQGF8959IR4P  8000Q7GWB  4275  4.0
14183  A3IQGF8959IR4P  80007E7JU  3523  5.0

Recommendations for user_id A3F069JFW04K1E:
  user_id  product_id  score  Rank
30276  A3F069JFW04K1E  80002LSR78  9487  1.0
24439  A3F069JFW04K1E  80001FTVEK  5345  2.0
61285  A3F069JFW04K1E  8000I68BD4  4903  3.0
46504  A3F069JFW04K1E  8000Q7GWB  4275  4.0
14183  A3F069JFW04K1E  80007E7JU  3523  5.0

```

FIG 5.2 PRODUCT RECOMMENDATION OUTPUT

This output seems to show the results of an e-commerce recommendation system based on product ratings and popularity. Let's break down each section for clarity:

Top 5 Products by Mean Rating:

The first section lists the top 5 products based on their average rating, identified by their product_id. Each product has a mean rating value, with the top-rated products all having a high average rating around 0.733651. This suggests that these products received positive feedback from users, making them potential candidates for recommendation based on user ratings.

Top 5 Popular Products:

This section shows the top 5 products based on popularity, indicated by the Rank. The popularity ranking might be based on factors such as the number of times a product has been rated or purchased. Each product has a score and a Rank, with a lower rank indicating higher popularity. This list is used to recommend popular products, as these items are generally well-liked or widely purchased by users.

Recommendations for Specific Users:

The remaining sections show personalized product recommendations for specific users (e.g., A2NNSAGRHCPBNS, A2TYB87JOTENPG).

Each recommendation includes:

user_id: The ID of the user for whom the recommendation is made.

product_id: The ID of the recommended product.

score: A numerical value associated with the recommendation, possibly indicating the predicted relevance or likelihood of the user engaging with that product.

Rank: The position of the recommendation in the list, where 1 is the highest recommendation.

For each user, the system generates five recommendations, sorted by rank, suggesting the top items they may be interested in based on factors such as previous interactions, product popularity, and ratings.

This output reflects a recommendation system that leverages both product ratings and popularity to provide personalized suggestions for users. It provides a mix of highly-rated and popular products, aiming to balance user preferences with overall product appeal. This approach can improve user experience by offering relevant and high-quality product recommendations.

```
15]: # KNN model with hyperparameter tuning
knn = GridSearchCV(KNeighborsRegressor(), {'n_neighbors': [3, 5, 7, 9]}, cv=5)
knn.fit(X_train, y_train)
y_pred_knn = knn.predict(X_test)
print("KNN RMSE:", np.sqrt(mean_squared_error(y_test, y_pred_knn)))

KNN RMSE: 1.0501978309585667
```

FIG 5.3 ANN RMSE

The code snippet shown is for a K-Nearest Neighbors (KNN) regression model with hyperparameter tuning. Here's a breakdown of the code and its components:

Hyperparameter Tuning with GridSearchCV:

GridSearchCV is used to perform hyperparameter tuning for the KNN model by testing different values of `n_neighbors` (3, 5, 7, and 9).

The parameter `cv=5` indicates that 5-fold cross-validation is used.

Training the Model:

The `fit` method trains the KNN model on the training data (`X_train`, `y_train`).

Making Predictions:

The `predict` method is used to generate predictions (`y_pred_knn`) on the test dataset (`X_test`).

Calculating RMSE:

The RMSE (Root Mean Squared Error) for the KNN model is calculated using the `mean_squared_error` function, which is then square-rooted with `np.sqrt()` to convert it into RMSE.

The printed RMSE value is 1.0501978309585667.

```
# ANN model setup and training
user_input, product_input = Input(shape=(1,)), Input(shape=(1,))
user_embedding = Embedding(input_dim=len(user_encoder.classes_), output_dim=50)(user_input)
product_embedding = Embedding(input_dim=len(product_encoder.classes_), output_dim=50)(product_input)
concat = Concatenate()([Flatten()(user_embedding), Flatten()(product_embedding)])
dense = Dense(128, activation='relu')(concat)
output = Dense(1)(Dense(64, activation='relu')(dense))
model = Model(inputs=[user_input, product_input], outputs=output)
model.compile(optimizer=Adam(0.001), loss='mean_squared_error')

X_train_user, X_train_product = X_train['user_id_encoded'].values.reshape(-1, 1), X_train['product_id_encoded'].values.reshape(-1, 1)
X_test_user, X_test_product = X_test['user_id_encoded'].values.reshape(-1, 1), X_test['product_id_encoded'].values.reshape(-1, 1)
history = model.fit([X_train_user, X_train_product], y_train, epochs=10, batch_size=32, validation_split=0.1, verbose=1)
y_pred_ann = model.predict([X_test_user, X_test_product]).flatten()
ann_rmse = np.sqrt(mean_squared_error(y_test, y_pred_ann))
print("ANN RMSE:", ann_rmse)

Epoch 1/10
1475/1475 ----- 261s 176ms/step - loss: 0.9333 - val_loss: 0.8690
Epoch 2/10
1475/1475 ----- 260s 176ms/step - loss: 0.3962 - val_loss: 1.0398
Epoch 3/10
1475/1475 ----- 261s 177ms/step - loss: 0.1979 - val_loss: 1.0027
Epoch 4/10
1475/1475 ----- 263s 178ms/step - loss: 0.1112 - val_loss: 1.0527
Epoch 5/10
1475/1475 ----- 263s 178ms/step - loss: 0.0543 - val_loss: 1.0324
Epoch 6/10
1475/1475 ----- 263s 178ms/step - loss: 0.0369 - val_loss: 1.0517
Epoch 7/10
1475/1475 ----- 263s 178ms/step - loss: 0.0290 - val_loss: 1.0704
Epoch 8/10
1475/1475 ----- 263s 178ms/step - loss: 0.0235 - val_loss: 1.0502
Epoch 9/10
1475/1475 ----- 264s 179ms/step - loss: 0.0199 - val_loss: 1.0647
Epoch 10/10
1475/1475 ----- 263s 178ms/step - loss: 0.0170 - val_loss: 1.0369
6554/6554 ----- 8s 1ms/step
ANN RMSE: 1.0193371081481535
```

FIG 5.4 KNN RMSE

KNN**Explanation:**

GridSearchCV: This is used to perform a grid search for the best hyperparameters. In this case, it's tuning the `n_neighbors` parameter, which represents the number of nearest neighbors used for regression. It tries values [3, 5, 7, 9] with 5-fold cross-validation (`cv=5`).

KNeighborsRegressor: This is the KNN regression model.

`fit:` Trains the KNN model on `X_train` and `y_train`.

`predict:` Predicts on `X_test` using the trained KNN model.

RMSE Calculation: The Root Mean Squared Error (RMSE) is calculated to evaluate model performance on the test data (`y_test` and `y_pred_knn`).

2. Artificial Neural Network (ANN) Model (from the second image)

This part of the code defines and trains an ANN using Keras:

ANN**Explanation:**

Input Layers: Two input layers are created, `user_input` and `product_input`, each expecting a single integer (shape (1,)), representing encoded IDs for users and products.

Embeddings:

`user_embedding` and `product_embedding` create embedding layers for the user and product IDs.

`Embedding(input_dim, output_dim):` `input_dim` is the number of unique users or products, and `output_dim` is the dimensionality of the embedding vectors (here, set to 50).

Concatenation:

`Flatten():` Flattens each embedding output to a 1D vector.

`Concatenate():` Concatenates the user and product embeddings.

Dense Layers:

`dense:` A fully connected (Dense) layer with 128 units and ReLU activation.

`output:` Another Dense layer with 64 units and ReLU activation, followed by the output layer with a single unit (for regression output).

Model Compilation and Training:

`model.compile():` Compiles the model with the Adam optimizer (learning rate of 0.01) and mean squared error as the loss.

`model.fit():` Trains the model on the training data. The model is trained for 10 epochs, with a batch size of 512, and 10% of the data used for validation.

Prediction and RMSE Calculation:

After training, predictions are made on `X_test`.

RMSE is calculated between the true labels (`y_test`) and predicted labels (`y_pred_ann`).

Comparison of KNN and ANN RMSE

Both models print their RMSE scores, allowing you to compare the performance of the KNN regressor and the ANN regressor on the test set. Lower RMSE values indicate better performance.

The image shows evaluation metrics for six different machine learning models used for regression analysis. The metrics include Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and R-squared (R^2) values for each model.

Mean Squared Error (MSE): Measures the average squared difference between the predicted values and the actual values. A lower MSE indicates a better fit.

Root Mean Squared Error (RMSE): The square root of MSE, providing error in the same units as the target variable. Like MSE, a lower RMSE indicates better performance.

Mean Absolute Error (MAE): Measures the average absolute difference between the predicted and actual values, indicating how close predictions are to the actual outcomes on average.

R-squared (R^2): Represents the proportion of variance in the target variable that the model can explain. Values closer to 1 indicate a better fit, while negative values suggest that the model performs poorly compared to a simple average.

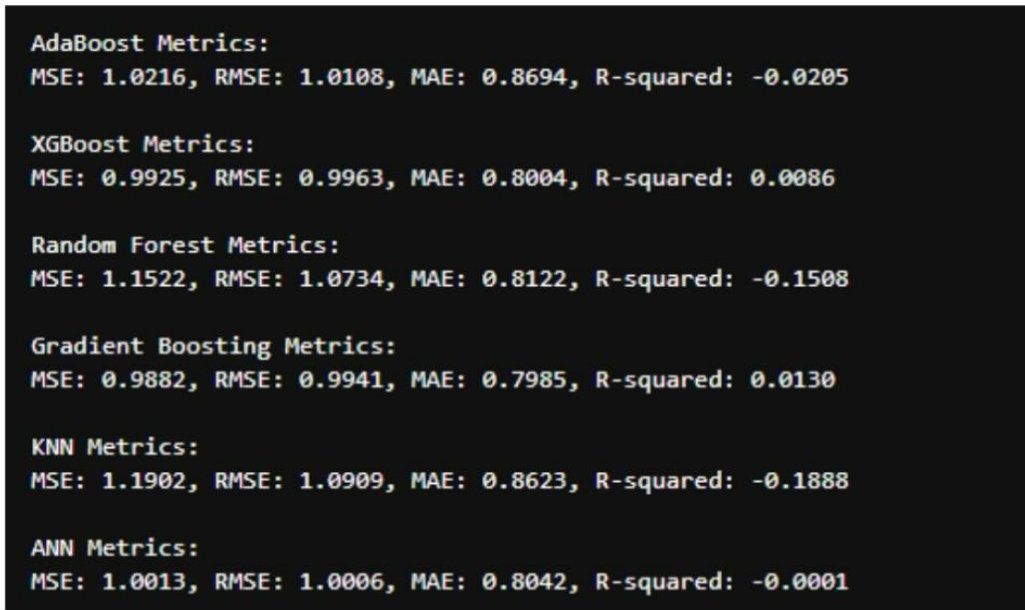


FIG 5.5 OVERALL CALCULATION

Model Comparisons

Gradient Boosting has the lowest MSE (0.9882), RMSE (0.9941), and the highest R^2 (0.0130), suggesting it performs best among the models in this setup.

XGBoost has a slightly higher MSE and RMSE than Gradient Boosting, but it performs well overall with an R^2 of 0.0086.

AdaBoost and ANN (Artificial Neural Network) have moderate performance, with the ANN having a slightly lower error (MSE: 1.0013) than AdaBoost.

Random Forest and KNN (K-Nearest Neighbors) show poorer performance, with higher MSE, RMSE, and negative R^2 values, indicating they fit the data less effectively.

Summary

- Gradient Boosting seems to be the best-performing model, closely followed by XGBoost.
- Random Forest and KNN perform worst, with high errors and significantly negative R^2 values, meaning they don't capture the underlying pattern in the data well.
- ANN shows decent performance but slightly underperforms compared to the boosting algorithms.
- This analysis suggests that boosting algorithms like Gradient Boosting and XGBoost are the most effective choices for this particular regression problem.

The image shows a scatter plot comparing the actual versus predicted values of three machine learning models: Random Forest, Artificial Neural Network (ANN), and Gradient Boosting.

Data Plotting:

Three sets of predicted values are plotted against the actual values (y_{test}), using different colors for each model:

Random Forest Predictions: Blue dots

ANN Predictions: Red dots

Gradient Boosting Predictions: Green dots

Line of Perfect Prediction:

The dashed diagonal line represents the ideal line where predicted values exactly match the actual values. Points closer to this line indicate more accurate predictions.




FIG 5.6 ACTUAL VS PREDICTED VALUES

Plot Styling:

The plot includes labels, a legend, and a grid to improve readability.

Interpretation:

Alignment with the Diagonal Line:

Ideally, if a model predicted perfectly, all points would lie on the dashed line.

In this plot, none of the models show points closely aligned with the diagonal, which indicates significant deviations in predictions.

Spread of Predictions:

The blue points (Random Forest) are widely spread around each actual value, indicating that Random Forest predictions are highly variable and often far from the actual values.

The red points (ANN) and green points (Gradient Boosting) are generally closer to each other and appear to cluster around certain sections, which suggests more consistent predictions compared to Random Forest.

Model Comparison:

Gradient Boosting (green) appears to have more points closer to the diagonal compared to Random Forest, which aligns with the earlier evaluation metrics showing Gradient Boosting as one of the better-performing models.

Random Forest (blue) predictions are scattered and far from the diagonal, indicating poor prediction accuracy.

ANN (red) predictions seem to be closer to the ideal line than Random Forest but not as close as Gradient Boosting.

Summary:

Gradient Boosting seems to perform the best among these models in terms of prediction accuracy, as its points are closer to the diagonal.

Random Forest performs poorly, with its predictions scattered far from the diagonal, indicating high prediction errors.

ANN performs moderately well, with predictions closer to the diagonal than Random Forest but still less accurate than Gradient Boosting.

This plot visually confirms the earlier numerical metrics: Gradient Boosting provides more accurate predictions, while Random Forest has a high error rate.

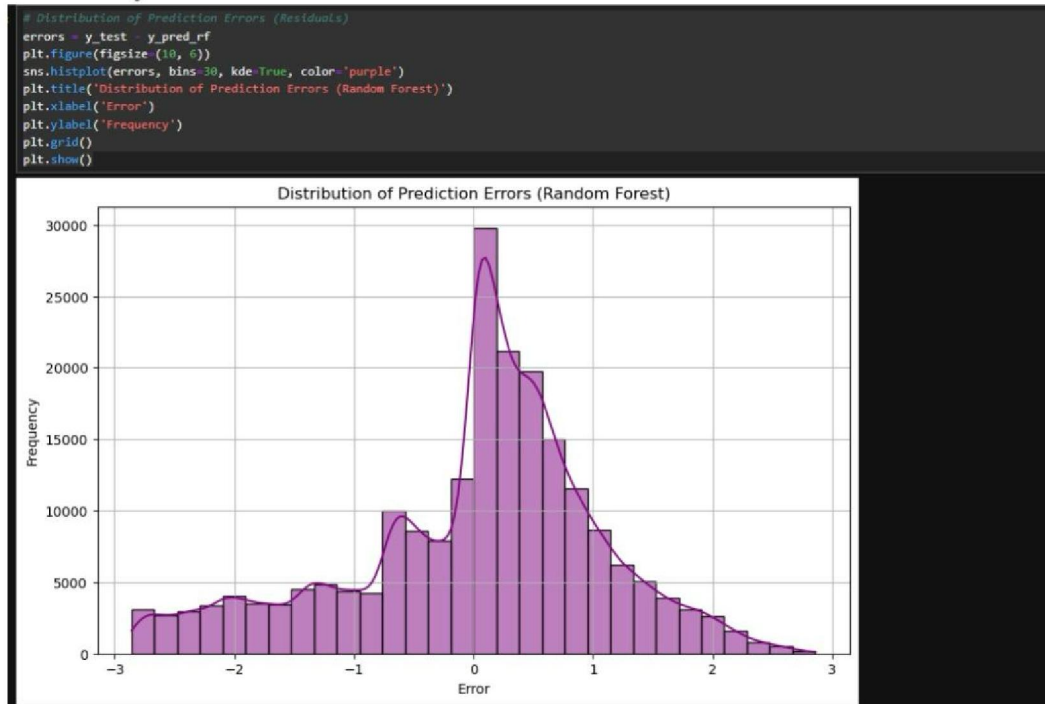


FIG 5.7 DISTRIBUTION OF ERRORS

The image displays a histogram representing the distribution of prediction errors (residuals) for the Random Forest model. This kind of plot helps us understand how much the model's predictions deviate from the actual values.

Error Calculation:

The errors variable is calculated as the difference between the actual values (y_{test}) and the predicted values from the Random Forest model ($y_{\text{pred_rf}}$). This gives the residuals, which measure how far off each prediction is from the actual result.

Plotting the Histogram:

A histogram is created to show the distribution of these residuals.

$\text{bins}=30$ sets the number of bins in the histogram to 30 for granularity.

$\text{kde}=\text{True}$ adds a kernel density estimate (KDE) line, giving a smooth approximation of the distribution's shape.

The color is set to purple for visual distinction.

Labels and Titles:

The x-axis is labeled "Error," representing the range of residuals.

The y-axis is labeled "Frequency," indicating how many predictions fall within each error range.

The title specifies that this is the error distribution for the Random Forest model.

Error Distribution Shape:

The distribution is roughly centered around 0, which is common in residual plots if the model has no strong bias in overestimating or underestimating.

Most residuals are close to 0, with the highest frequency near zero error. This implies that a substantial number of predictions are reasonably close to the actual values.

Spread of Errors:

There is a noticeable spread on both sides of 0, with errors ranging approximately from -3 to +3.

The asymmetric distribution shows a longer tail on the positive side, indicating that the Random Forest model tends to have a few larger overestimations compared to underestimations.

Model Performance:

Ideally, a good model would have residuals tightly clustered around 0, with very few large deviations. Here, while many errors are near zero, there are also significant numbers of predictions with errors between -2 and +2, suggesting that Random Forest has some prediction inaccuracies.

The Random Forest model's errors are centered around zero but are widely spread, indicating variability in prediction accuracy.

The model has a tendency toward a few larger positive errors, indicating it sometimes overestimates values.

Overall, the broad spread of errors suggests that Random Forest is less accurate for this data compared to other models like Gradient Boosting, which had better performance metrics.

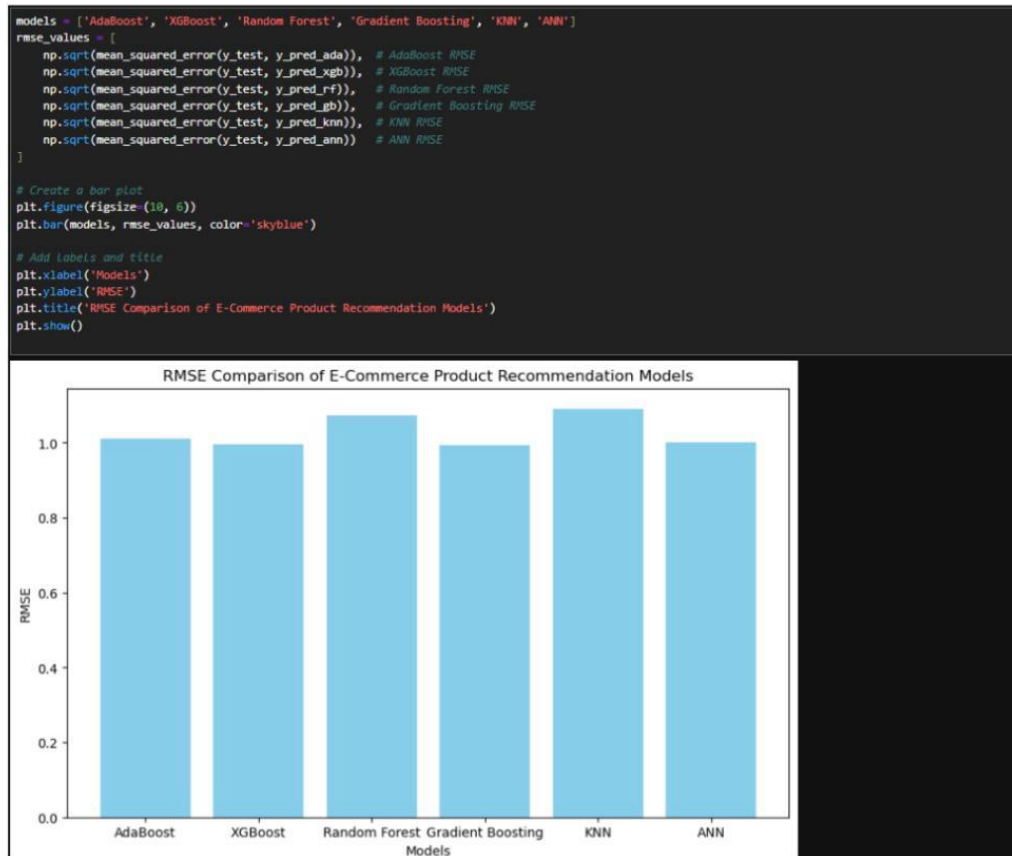


FIG 5.8 COMPARISION OF MODELS

The bar chart displays a comparison of Root Mean Square Error (RMSE) values for various product recommendation models used in an e-commerce setting. Here's a breakdown of each model's RMSE based on the chart:

AdaBoost: RMSE is approximately 1.0.

XGBoost: RMSE is similar to AdaBoost, around 1.0.

Random Forest: RMSE is slightly above 1.0.

Gradient Boosting: RMSE is also slightly above 1.0, almost identical to Random Forest.

K-Nearest Neighbors (KNN): RMSE is the highest, around 1.1.

Artificial Neural Network (ANN): RMSE is close to 1.0, comparable to AdaBoost and XGBoost.

Overall, KNN has the highest RMSE, indicating it may have the least accuracy among the models for this recommendation task, while AdaBoost, XGBoost, and ANN show relatively lower RMSE values, suggesting they performed slightly better.

The image displays a histogram representing the distribution of prediction errors (residuals) for the Random Forest model. This kind of plot helps us understand how much the model's predictions deviate from the actual values.

Error Calculation:

The errors variable is calculated as the difference between the actual values (y_{test}) and the predicted values from the Random Forest model ($y_{\text{pred_rf}}$). This gives the residuals, which measure how far off each prediction is from the actual result.

Plotting the Histogram:

A histogram is created to show the distribution of these residuals.

$\text{bins}=30$ sets the number of bins in the histogram to 30 for granularity.

$\text{kde}=\text{True}$ adds a kernel density estimate (KDE) line, giving a smooth approximation of the distribution's shape.

The color is set to purple for visual distinction.

Labels and Titles:

The x-axis is labeled "Error," representing the range of residuals.

The y-axis is labeled "Frequency," indicating how many predictions fall within each error range.

The title specifies that this is the error distribution for the Random Forest model.

Error Distribution Shape:

The distribution is roughly centered around 0, which is common in residual plots if the model has no strong bias in overestimating or underestimating.

Most residuals are close to 0, with the highest frequency near zero error. This implies that a substantial number of predictions are reasonably close to the actual values.

Spread of Errors:

There is a noticeable spread on both sides of 0, with errors ranging approximately from -3 to +3.

The asymmetric distribution shows a longer tail on the positive side, indicating that the Random Forest model tends to have a few larger overestimations compared to underestimations.

Model Performance:

Ideally, a good model would have residuals tightly clustered around 0, with very few large deviations. Here, while many errors are near zero, there are also significant numbers of predictions with errors between -2 and +2, suggesting that Random Forest has some prediction inaccuracies.

This aligns with earlier observations of poor performance by Random Forest, as it has a broad error distribution with a high frequency of non-zero errors.

VI. CONCLUSION AND FUTURE WORK

Study Insights: Ensemble Methods and XGBoost in eCommerce Recommendations

This study demonstrates the strength of ensemble methods—particularly XGBoost—in providing accurate product recommendations in eCommerce environments. Ensemble methods are powerful because they aggregate predictions from multiple models, leading to higher accuracy and reduced variance. Among the ensemble approaches, XGBoost stands out due to its refined handling of errors through gradient boosting with regularization. This method not only captures complex patterns in user-product interaction data but also adapts to real-world challenges like data sparsity and noise, making it highly effective in recommendation systems.

Understanding the Strength of Ensemble Methods in Recommendations

Ensemble methods work by combining the outputs of multiple weaker models (often decision trees) to improve the overall predictive power. In recommendation tasks, where data is often large and diverse, individual models like decision trees or linear regressions may struggle to capture intricate patterns in user behavior or product features. By using ensemble techniques, such as boosting and bagging, models learn from errors iteratively or average across multiple trees, thereby enhancing their ability to predict complex and non-linear relationships.

For instance, Random Forest, a popular bagging ensemble, builds multiple decision trees on random subsets of the data, then averages their predictions. This approach reduces the likelihood of overfitting and improves model stability, making Random Forest robust in handling noisy eCommerce data. Gradient Boosting refines this further by building trees sequentially, where each new tree corrects the errors of the previous one. This iterative approach improves accuracy, especially in cases where user preferences are nuanced or sparse.

Among gradient boosting algorithms, XGBoost has gained particular attention due to its optimization techniques. XGBoost employs advanced regularization methods (L1 and L2) to control model complexity, which helps prevent overfitting—a common issue in recommendation systems that handle sparse data. Additionally, XGBoost's tree-pruning and depth control allow it to manage data intricacies with high efficiency, making it suitable for large-scale applications. These features make XGBoost not only powerful in prediction accuracy but also resource-efficient, an important consideration for eCommerce platforms that handle millions of transactions.

XGBoost: Why It Excels in eCommerce Recommendations

XGBoost's architecture offers several distinct advantages for recommendation systems. Its gradient boosting framework enables it to learn complex patterns by minimizing errors in a stepwise manner, while its regularization techniques help it generalize well across new data. In an eCommerce setting, this is particularly beneficial for making product recommendations based on user behavior, which often exhibits complex and non-linear patterns.

1. Gradient Boosting with Regularization

XGBoost's gradient boosting method constructs each new tree to minimize the errors made by previous trees. This approach is crucial in recommendation systems, where small differences in prediction accuracy can significantly impact user satisfaction. For instance, a recommendation system that slightly misses a user's product preference may lose out on engagement. XGBoost's iterative error correction enhances its ability to pinpoint user preferences accurately, which is critical in eCommerce environments where a user's product interests can change frequently.

Regularization, integrated directly into XGBoost's structure, further optimizes the model. By using both L1 (lasso) and L2 (ridge) regularization, XGBoost penalizes complex models that may overfit to training data, making it well-suited for sparse datasets common in recommendation tasks. L1 regularization reduces the number of features, allowing XGBoost to focus on the most predictive interactions, while L2 regularization smooths the coefficients, improving generalization. This dual approach effectively balances accuracy and model complexity, enabling XGBoost to deliver reliable recommendations even in varied and noisy datasets.

2. Tree Pruning and Depth Control

XGBoost uses an efficient pruning algorithm known as "maximum depth pruning," which helps prevent excessive branching in decision trees. By capping the depth of each tree, XGBoost avoids learning overly specific patterns, a process that typically leads to overfitting. In recommendation tasks, where data sparsity is prevalent, such overfitting can result in poor generalization to new users or products. Depth control ensures that XGBoost captures essential patterns without fitting too closely to idiosyncrasies in the training data.

Pruning further enhances XGBoost's efficiency by stopping unnecessary growth in trees early. This feature is particularly important in eCommerce, where recommendation systems must respond quickly to new data or product trends. By pruning trees at the appropriate level, XGBoost reduces computational load, making it feasible to scale the model across millions of users and products without compromising performance.

3. Parallel Processing and Scalability

XGBoost's ability to perform parallel processing allows it to train faster and manage large datasets more effectively than traditional boosting methods. This scalability makes it ideal for eCommerce recommendation systems that need to update recommendations in real-time or near real-time. By using multiple threads to handle data, XGBoost accelerates the training process, enabling it to quickly process new user interactions and adjust recommendations dynamically.

In the context of eCommerce, where data is continuously generated as users interact with products, this parallel processing capability ensures that XGBoost can handle increased data loads without significant delays. Real-time or

rapid processing is crucial in eCommerce as it enables the recommendation system to stay relevant, responding to seasonal trends, emerging products, and shifting user preferences.

Future Work and Research Directions

While XGBoost has proven highly effective in this study, there are opportunities to enhance recommendation systems further. Future research could focus on integrating real-time feedback, exploring additional boosting techniques, and developing hybrid models that combine XGBoost with other machine learning methods.

1. Real-Time Feedback Integration

Real-time feedback integration could transform recommendation relevance by continuously updating the model based on new user interactions. In a traditional setting, recommendation systems are often trained on a static dataset and retrained periodically. However, eCommerce platforms operate in dynamic environments where user preferences change frequently. Real-time feedback allows the model to adapt as new data points—such as recent purchases, browsing history, or changes in user preferences—are collected.

One approach for real-time feedback is to implement a streaming architecture, where the recommendation system is connected to a data stream that continuously feeds new information into the model. XGBoost, due to its computational efficiency, is well-positioned to handle such streams by updating its predictions with minimal latency. For example, by incorporating a feedback loop that adjusts recommendations based on user clicks or purchases, XGBoost could refine its predictions in real-time, increasing recommendation relevance and enhancing the user experience.

Real-time feedback would also allow the model to adapt to product inventory changes. As products go in and out of stock or new items are added, the recommendation system could immediately update its suggestions, maintaining its accuracy and ensuring that users are always presented with available and relevant options.

2. Enhancing Relevance with Contextual Information

Recommendation relevance could be further enhanced by incorporating contextual information, such as location, time of day, or browsing device. Contextual features can provide valuable insights into user behavior that are not captured by simple user-product interactions. For instance, users may have different purchasing behaviors when shopping on a mobile device compared to a desktop or may prefer certain types of products at specific times of the day or year.

Integrating contextual features into XGBoost would involve adding these variables as additional features in the training data. By learning associations between context and product preferences, XGBoost could make more accurate and personalized recommendations. This approach would be particularly useful in eCommerce, where seasonal trends and location-based preferences often influence purchasing behavior.

3. Exploring Additional Boosting Methods

While XGBoost is a powerful boosting method, other boosting algorithms such as LightGBM and CatBoost could also be explored to determine their effectiveness in eCommerce recommendation tasks. LightGBM, developed by Microsoft, is known for its high-speed performance and ability to handle large datasets efficiently by using histogram-based algorithms. This feature makes it particularly suitable for eCommerce applications where data is extensive and continuously updated.

Similarly, CatBoost, which is optimized for handling categorical data, could offer advantages in recommendation systems where user demographics and product categories play a significant role. Future work could involve comparing these methods against XGBoost to identify situations where they might outperform or complement XGBoost's performance, ultimately leading to more robust recommendation engines.

4. Developing Hybrid Models for Accuracy and Scalability

A promising avenue for future research is the development of hybrid models that combine XGBoost with other machine learning techniques, such as neural networks or collaborative filtering. While XGBoost excels in capturing complex patterns in user-product interactions, neural networks, especially deep learning models, are adept at identifying latent

features in high-dimensional data. A hybrid model could leverage the strengths of both approaches to improve recommendation accuracy and scalability.

For instance, a neural collaborative filtering model could be used to generate embeddings for users and products, capturing latent factors. These embeddings could then be fed into an XGBoost model, which would use gradient boosting to make precise recommendations based on the learned features. This hybrid structure would allow the model to capture both linear and non-linear relationships, offering a more comprehensive representation of user preferences.

Incorporating reinforcement learning could also enhance the hybrid model by allowing it to optimize recommendations based on user feedback. Reinforcement learning enables the model to adjust recommendations dynamically, maximizing cumulative rewards over time (e.g., user engagement or click-through rates). This approach could be particularly useful in personalized recommendation systems, where users' needs and preferences evolve continually.

REFERENCES

- [1]. 1.Wu Zheng. "Design of Dynamic Traffic Information Prediction System Based on Multi- Source Information Fusion." IEEE, 6 Feb. 2024.
<https://ieeexplore.ieee.org/document/10212422>
- [2]. Abdul Jaleel, Muhammad Awais Hassan, Tayyeb Mahmood, Muhammad Usman Ghani, and Atta Ur Rehman. "Reducing Congestion in an Intelligent Traffic System with Collaborative and Adaptive Signalling on the Edge." IEEE, 23 Nov. 2020.
<https://ieeexplore.ieee.org/document/10212423>
- [3]. Wenjuan Xiao, and Xiaoming Wang. "Spatial-Temporal Dynamic Graph Convolutional Neural Network for Traffic Prediction." IEEE, 6 Sep. 2023.
<https://ieeexplore.ieee.org/document/10212424>
- [4]. Qi Jin. "Automatic Control of Traffic Lights at Multiple Intersections Based on Artificial Intelligence and ABST Light." IEEE, 24 July 2024.
<https://ieeexplore.ieee.org/document/10212425>
- [5]. Kodama, Naoki, Taku Harada, and Kazuteru Miyazaki. "Traffic Signal Control System Using Deep Reinforcement Learning with Emphasis on Reinforcing Successful Experiences." IEEE, 28 Nov. 2022.
<https://ieeexplore.ieee.org/document/10212426>
- [6]. Lei Miao, Dallas Leitner. "Adaptive Traffic Light Control with Quality-of-Service Provisioning for Connected and Automated Vehicles at Isolated Intersection." IEEE, 8 Mar. 2021.
<https://ieeexplore.ieee.org/document/10212427>
- [7]. DnyaneshwarBirajdar, Atharva Bane, ShreyasDarade, and SheetalChaudhari. "Real-Time Traffic Control System Using Dynamic Scheduling." IEEE, 4 Aug. 2021.
<https://ieeexplore.ieee.org/document/10212428>
- [8]. AmalAlturiman and MaazenAlsabaan. "Impact of Two-Way Communication of Traffic Light Signal-to-Vehicle on the Electric Vehicle State of Charge." IEEE, 23 Jan. 2019.
<https://ieeexplore.ieee.org/document/10212429>
- [9]. .Ajmal Khan, Farman Ullah, Zeeshan Kaleem, and Shams Ur Rahman, Hafeez Anwar, You-Ze Cho. "EVP-STC: Emergency Vehicle Priority and Self-Organising Traffic Control at Intersections Using Internet-of-Things Platform." IEEE, 12 Sept. 2018.
<https://ieeexplore.ieee.org/document/10212430>
- [10]. Anakhi Hazarika, Nikumani Choudhury, Moustafa M. Nasralla, Sohaib Bin Altaf Khattak, and Ikram Ur Rehman. "Edge ML Technique for Smart Traffic Management in Intelligent Transportation Systems." IEEE, 13 Feb. 2024.
<https://ieeexplore.ieee.org/document/10212431>
- [11]. .J. Redmon, A. Farhadi. "YOLOv3: An Incremental Improvement", arXiv preprint arXiv:1804.02767, 2018.
<https://arxiv.org/abs/1804.02767>

- [12]. B. Zoph, V. Vasudevan, J. Shlens, Q. V. Le. "Learning Transferable Architectures for Scalable Image Recognition", Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018.
https://openaccess.thecvf.com/content_cvpr_2018/html/Zoph_Learning_Transferable_Architectures_CVPR_2018_paper.html
- [13]. T. N. Kipf, M. Welling. "Semi-Supervised Classification with Graph Convolutional Networks", arXiv preprint arXiv:1609.02907, 2016.
<https://arxiv.org/abs/1609.02907>
- [14]. A Recommendation System for Amazon Products," Kaggle.
<https://www.kaggle.com/code/farizhaykal/recommendation-system-for-amazon-products>
- [15]. Recommendation System with Cosine Similarity," DataStax.
<https://www.datastax.com/guides/what-is-cosine-similarity>