

Optimizing Recommendation Performance with a Multi-Stage k-Means, DNN, RBM, and k-NN Pipeline

Ayush Yajnik¹, Vishal Sharma²

Jersey City, New Jersey, USA
Email: [ayushyajnik1\[at\]outlook.com](mailto:ayushyajnik1[at]outlook.com)

New Delhi, India
Email: [data.with.vishal\[at\]gmail.com](mailto:data.with.vishal[at]gmail.com)

Abstract: Recommendation systems help users navigate vast item catalogs, yet traditional collaborative- and content-based filtering suffer from data-sparsity, cold-start issues, and limited ability to model complex user item relationships. To overcome these challenges, we present a unified hybrid pipeline that first partitions the item space with k-means clustering, then employs a deep neural network for feature extraction and cluster selection, refines selections through a Restricted Boltzmann Machine, and finally delivers item suggestions using k-Nearest Neighbors. Experiments on the Kaggle Spotify dataset after z-score normalization and SMOTE-based class-balancing show that our deep network attains an average F1-score of 0.97, RBM refinement boosts within-cluster accuracy, and the final k-NN stage yields superior Precision, Recall, NDCG, and MAP compared with baseline collaborative-filtering and matrix-factorization models. These results demonstrate that orchestrating complementary algorithms in a multi-stage workflow produces robust, scalable, and highly accurate recommendations suitable for real-world deployment.

Keywords: Hybrid recommendation, k-means clustering, deep neural network, Restricted Boltzmann Machine, k-Nearest Neighbors, Spotify dataset, SMOTE, evaluation metrics

1. Introduction

Recommendation systems have become indispensable for digital platforms, powering personalized experiences across domains such as e-commerce, streaming services, and social media. These systems assist users in discovering relevant items from extensive catalogs, thereby enhancing user satisfaction and platform engagement. Traditional recommendation approaches, including collaborative filtering [17] and content-based filtering [18], have been widely adopted due to their simplicity and effectiveness. [16] However, these methods often struggle with challenges such as data sparsity, the cold start problem, and an inability to capture complex user-item relationships.

Collaborative filtering, one of the most popular recommendation techniques, relies on user-item interaction data to infer preferences [16]. However, its effectiveness diminishes in cases of sparse data or new users and items with no prior interactions, known as the cold start problem. Content-based filtering, on the other hand, utilizes item attributes to recommend similar items to those a user has engaged with [18]. While this approach alleviates some limitations of collaborative filtering, it often fails to generalize beyond the user's existing preferences, leading to limited diversity in recommendations.

Hybrid recommendation systems, which combine multiple techniques, have emerged as a promising solution to address these challenges. For example, matrix factorization techniques, such as Singular Value Decomposition (SVD), paired with deep learning models, have demonstrated improved accuracy in capturing latent user-item interactions [19][20]. Clustering-based methods, like k-means, have also been employed to group items or users into clusters, enhancing scalability and reducing computational

complexity. Despite these advances, many existing hybrid models focus on integrating only a few components, leaving untapped potential for combining diverse algorithms in a cohesive pipeline.

In this paper, we propose a novel hybrid recommendation system that integrates k-means clustering, deep neural networks (DNNs), Restricted Boltzmann Machines (RBMs), and k-Nearest Neighbors (k-NN) into a multi-stage pipeline. Our approach leverages the unique strengths of each algorithm to address the limitations of traditional and hybrid systems. K-means clustering is employed to partition the item space into cohesive groups, reducing the search space and improving computational efficiency. DNNs are used to perform feature extraction and classification, achieving high accuracy in selecting the most relevant clusters. RBMs further refine the selection process by probabilistically modeling intricate dependencies within the chosen clusters. Finally, k-NN generates precise and personalized recommendations by identifying similar items within the refined cluster.

To evaluate the effectiveness of our system, we use the Spotify dataset from Kaggle, which contains a diverse collection of music tracks enriched with audio features. Addressing class imbalance, we employ the Synthetic Minority Over-sampling Technique (SMOTE) [21] to balance the data distribution before training. Data normalization is also applied to ensure effective convergence of the learning algorithms. Our results demonstrate significant improvements in key performance metrics, including precision, recall, F1-score, Normalized Discounted Cumulative Gain (NDCG), and Mean Average Precision (MAP), compared to baseline methods such as collaborative filtering and matrix factorization.

This paper contributes to the field by demonstrating the value of combining diverse algorithms in a structured, multi-stage

pipeline. The proposed system not only achieves high recommendation accuracy but also addresses critical challenges like data sparsity and cold start, paving the way for robust and scalable recommendation systems in various domains.

1.1 K-Means Clustering

K-means clustering is one of the simplest yet most widely used unsupervised learning algorithms for grouping similar data points into clusters. The algorithm divides a set of n data points into k non-overlapping clusters, where each data point belongs to the cluster with the nearest centroid. The basic goal is to minimize the within-cluster sum of squares (WCSS), which is the sum of squared Euclidean distances between each point and its cluster's centroid.

Mathematical Formulation:

The K-means algorithm works by iteratively optimizing an objective function that measures the compactness of clusters. The objective function is:

$$J = \sum_{i=1}^k \sum_{x \in C_i} \|x - \mu_i\|^2,$$

where:

- k is the number of clusters,
- C_i is the set of points assigned to cluster i ,
- μ_i is the centroid of cluster i ,
- x is a data point.

The K-means algorithm proceeds as follows:

- 1) Initialization: Randomly initialize k centroids.
- 2) Assignment Step: Assign each data point to the nearest centroid, forming clusters.
- 3) Update Step: Recompute the centroids of each cluster based on the mean of the points in the cluster.
- 4) Convergence: Repeat steps 2 and 3 until the centroids no longer change significantly, signaling convergence.

While K-means is computationally efficient and simple to implement, it has some limitations:

- It requires specifying k in advance.
- It can get stuck in local minima depending on the initial centroid positions.
- It assumes clusters to be spherical and of equal size, which may not always be the case.

In recommendation systems, K-means is often used for preprocessing to reduce the dimensionality and complexity of the data. By grouping users or items into clusters, it becomes easier to perform subsequent tasks like predicting ratings or suggesting new items. For example, grouping users with similar preferences can help in collaborative filtering by identifying shared preferences between clusters.

K-means clustering has been extensively used in recommendation systems for various preprocessing and grouping tasks:

- 1) **User Segmentation:** By clustering users based on their preferences or behavioral data (e.g., purchase history, clickstream data), K-means helps in creating segments of

users with similar interests. This segmentation aids in personalized recommendations for groups of users rather than individuals, improving scalability [17].

- 2) **Item Clustering:** Similarly, K-means can cluster items based on their features (e.g., genre, price range, or attributes) to recommend similar items to users. For example, if a user interacts with an item from a specific cluster, other items from the same cluster can be suggested.
- 3) **Cold-Start Problem:** In situations where new users or items lack sufficient interaction data, clustering can help provide recommendations based on their similarity to existing clusters. New users can be assigned to clusters based on demographic or initial interaction data, and recommendations can be drawn from the cluster's common preferences.
- 4) **Dimensionality Reduction:** High-dimensional datasets in recommendation systems (e.g., user-item interaction matrices) can be simplified by grouping similar users or items into clusters. This reduces the computational complexity of collaborative filtering methods, as predictions can be made at the cluster level instead of the individual level.

1.2 Deep Neural Network (DNNs)

Deep Neural Networks (DNNs) have revolutionized machine learning and recommendation systems by enabling the extraction of complex features and modeling intricate relationships between users and items. A DNN consists of multiple layers of interconnected neurons that can learn non-linear mappings from inputs to outputs. In recommendation systems, DNNs are often used for predicting user preferences by capturing latent features through a process of hierarchical learning.

The loss function used in a typical DNN is cross-entropy loss, which measures the difference between the true and predicted class probabilities. The formula for cross-entropy loss is:

$$L = - \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C y_{ij} \log(\hat{y}_{ij}),$$

where:

- N is the number of samples,
- C is the number of classes (clusters or ratings),
- y_{ij} is the true label,
- \hat{y}_{ij} is the predicted probability.

The training of a DNN involves the backpropagation algorithm, which computes gradients and updates the weights to minimize the loss function. DNNs can learn complex user-item interactions, making them effective in collaborative filtering, where the goal is to predict user preferences based on historical interactions.

One common application of DNNs in recommendation systems is Neural Collaborative Filtering (NCF), where a DNN is used to model the interaction between users and items. In this framework, the user and item embeddings are learned through the network layers, allowing the model to capture intricate user-item relationships that traditional methods may miss. Key applications include:

- 1) **Neural Collaborative Filtering (NCF):** NCF frameworks, such as Neural Matrix Factorization (NeuMF), utilize DNNs to model user-item interactions. Instead of relying solely on traditional latent factor models, NCF combines user and item embeddings with non-linear transformations in a neural network to capture complex and non-linear relationships [20].
- 2) **Content-Based Recommendations:** DNNs are widely used in content-based recommendation systems for learning representations from high-dimensional data like text, images, and audio. For instance, Convolutional Neural Networks (CNNs) can extract features from images for item recommendations, while Recurrent Neural Networks (RNNs) can be applied to sequence-based recommendations, such as predicting the next item a user will interact with.
- 2) **Hybrid Systems:** RBMs are often integrated into hybrid recommendation systems to complement other algorithms like matrix factorization. Their generative nature helps to learn complex distributions in user preferences, which can be used alongside discriminative methods for better predictions. [25]
- 3) **Dimensionality Reduction for Sparse Data:** RBMs are effective for handling sparse and high-dimensional data, making them a good fit for recommendation systems where user-item interaction matrices are often incomplete and sparse. RBMs reduce the dimensionality of the data while preserving important patterns, improving computational efficiency and accuracy. [26]

1.3 Restricted Boltzman Machines (RBMs):

Restricted Boltzmann Machines (RBMs) are generative stochastic neural networks used for dimensionality reduction, feature learning, and collaborative filtering in recommendation systems. An RBM consists of two layers: a visible layer v that represents the observed data, and a hidden layer h that represents the latent features. The layers are fully connected, but there are no connections within each layer, making it "restricted."

The energy function for an RBM is defined as:

$$E(v, h) = - \sum_i v_i b_i - \sum_j h_j c_j - \sum_{i,j} v_i h_j w_{ij},$$

where:

- v_i and h_j represent the visible and hidden layer units,
- b_i and c_j are the biases for the visible and hidden layers,
- w_{ij} represents the weights between visible and hidden units.

The probability of a visible vector v is given by:

$$P(v) = \frac{1}{Z} \sum_h e^{-E(v,h)},$$

where Z is the partition function that ensures the probabilities sum to one.

RBMs are trained using contrastive divergence, which approximates the gradients for updating weights. The hidden layer learns to capture dependencies between visible variables, making RBMs particularly useful in collaborative filtering, where the visible layer represents user-item interactions, and the hidden layer captures latent features such as user preferences.

In recommendation systems, RBMs can be used to model user-item interactions and recommend items based on learned patterns. They are often used in hybrid systems where their generative nature complements other algorithms like matrix factorization or collaborative filtering. Some of their notable applications in recommendation systems include:

- 1) **Collaborative Filtering:** RBMs have been successfully used in collaborative filtering tasks, where the visible layer represents user-item interaction matrices, and the hidden layer captures latent features. By reconstructing the input data, RBMs can predict missing values, such as unobserved user ratings. [24]

1.4 k-Nearest Neighbors (k-NNs)

k-Nearest Neighbors (k-NN) is a simple and intuitive algorithm used for classification and regression tasks. In the context of recommendation systems, k-NN is used to identify items or users that are similar based on a distance metric, typically Euclidean distance. The k-NN algorithm works by finding the k nearest neighbors to a given data point and using the neighbors' characteristics to predict the target value.

The distance between two points x and y in a d -dimensional space is computed as:

$$\text{distance}(x, y) = \sqrt{\sum_{i=1}^d (x_i - y_i)^2},$$

where x_i and y_i are the features of the points x and y , respectively.

In recommendation systems, k-NN can be used for both user-based and item-based collaborative filtering. In user-based filtering, recommendations are made by finding users who are similar to the target user and suggesting items that they have rated highly. In item-based filtering, recommendations are made by finding items similar to the target item and suggesting those to the user.

1.5 Hybrid approach

The integration of multiple algorithms into a hybrid system has proven effective in overcoming the limitations of individual techniques. Hybrid approaches combine different algorithms to leverage their complementary strengths. For example, combining collaborative filtering with content-based methods can help address the cold-start problem, where there is insufficient user-item interaction data.

In recommendation systems, hybrid methods can take various forms:

- **Weighted Hybrid:** Different algorithms are combined by giving each one a weight. The final recommendation is based on a weighted average of the predictions from each algorithm.
- **Switching Hybrid:** The system switches between different algorithms depending on the context. For example, a content-based approach may be used for new items, while collaborative filtering is used for mature items with sufficient interaction data.
- **Cascade Hybrid:** One algorithm is used to filter out irrelevant items, and another algorithm is used for ranking the remaining items.

The remainder of this paper is organized as follows. **Section 2** surveys the relevant literature on traditional, deep-learning, and hybrid recommendation techniques. **Section 3** details our proposed multi-stage pipeline, including data preprocessing (3.1), k -means clustering (3.2), deep-neural-network cluster selection (3.3), Restricted Boltzmann Machine refinement (3.4), and k -Nearest-Neighbor recommendation generation (3.5). **Section 4** describes the experimental setup, reports the results on the Kaggle Spotify dataset, and discusses comparative performance against baseline models. Finally, **Section 5** concludes the study and outlines avenues for future research.

2. Related Work

2.1 Clustering-Based Recommendation Approaches

Partitioning users or items into coherent groups is a proven way to cut computation and ease sparsity. **Das et al.** introduced a scalable clustering framework that reduced collaborative-filtering (CF) runtime by up to 90 % on MovieLens-20M while keeping accuracy stable [1]. Building on that idea, **Zhang et al.** proposed KUR-CF, an enhanced K -means algorithm that weights user-attribute variance before similarity calculation; KUR-CF improved precision and recall by 60 % and 35 %, respectively, over vanilla CF on MovieLens-1M [2]. To tackle cold-start users and items, **Panyatip et al.** combined fuzzy c-means (FCM) clustering with item-based k -nearest neighbors (k -NN), achieving ≈ 85 % precision for new users and 87 % for new items [3]. A graph-driven alternative, **Darban and Valipour's GHRS**, first extracts auto-encoded features, then clusters users to boost cold-start accuracy and outperforms standard hybrids on the same benchmark [4]. The main strength of clustering is lower time/space cost and better sparsity handling, though cluster granularity and boundary overlap remain open issues.

2.2 Neighborhood-Based Collaborative Filtering (k -NN)

Although k -NN CF is simple and explainable, it falters with sparse data and new entities. **Behera et al.** mitigated sparsity by weighting a content-based KNN with an RBM predictor, yielding notably higher precision than either method alone [5]. To alleviate the new-user problem, **Latrech et al.** appended a deep demographic-filtering stage (CoDFi-DL) ahead of neighborhood CF, cutting RMSE to 0.571 on ML-1M [6]. **Nguyen et al.** made similarity adaptive by injecting user-cognition factors after initial clustering, which lifted MAP and NDCG versus baseline KNN [7]. Finally, **Li et al.** filled missing ratings with KNN, then trained XGBoost for the final prediction; this two-step hybrid lowered MAE and avoided local optima [8]. In short, modern work treats k -NN as a lightweight module inside richer hybrids to balance transparency with accuracy.

2.2 Deep-Learning and Hybrid Models

Deep neural networks (DNNs) learn non-linear interactions and latent factors that classic CF cannot capture. A recent survey by **Gheewala et al.** confirms that DNNs now dominate state-of-the-art ranking and prediction tasks, especially when fused with other techniques [9]. **Sami et al.** embedded a multilayer NCF block, a sequence-aware RNN, and content-

based filtering into a single hybrid (HRS-IU-DL); it achieved the best RMSE and F1 on MovieLens-100K among all contenders [10]. Addressing sparsity and cold start, **Heidari et al.** introduced ADLRS, an attention-based model that merges review-text embeddings with matrix factorization, outperforming MF on sparse sets [11]. For rich content scenarios, **Chetana et al.** used an improved DenseNet to mine movie metadata; their class-balanced loss trimmed RMSE by 5–10 % and raised F1 by ~ 2 % [12]. A neural-plus-traditional blend by **Sharma and Shakya** combined SVD++, self-organizing clustering, and an ANN, yielding higher recall than single-model baselines [13]. These studies prove that deep components boost accuracy and diversity, but introduce training cost and interpretability hurdles.

2.3 RBM-Based Collaborative Filtering

Restricted Boltzmann Machines (RBMs) still attract interest when enriched with side information. The hybrid of **Behera et al.** (content-KNN + RBM) illustrates how probabilistic modeling plus content features beat pure RBM or KNN [5]. **Wu et al.** coupled a Gaussian RBM with a CNN that embeds review text; the merged visible layer improved rating-prediction accuracy and demonstrated feasibility for text-aware CF [14]. While RBMs capture binary/implicit feedback well, they are training-intensive and sensitive to hyper-parameters; therefore, most recent systems deploy them only as sub-modules inside broader hybrids.

2.4 Summary and Challenges

Recent literature (2021–2025) converges on **hybrid recommender systems** that orchestrate clustering, neighborhood CF, deep learning, and probabilistic models to tackle data sparsity, cold start, scalability, and accuracy in unison. Clustering cuts complexity; k -NN offers transparency; DNNs and RBMs learn rich latent factors; their combination consistently outperforms individual methods. Remaining issues include model interpretability, hyper-parameter tuning, and the engineering overhead of maintaining multi-component pipelines. Nevertheless, systematic reviews confirm hybrids as the most promising path forward [15].

3. Methodology

This section elaborates on the proposed hybrid model, the data preprocessing steps, the architecture of the pipeline, and the mathematical formulations underpinning the methodology. The hybrid approach integrates k -Means clustering, Deep Neural Networks (DNNs), Restricted Boltzmann Machines (RBMs), and k -Nearest Neighbors (k -NN) into a cohesive framework to enhance recommendation accuracy and scalability.

3.1 Hybrid Model Architecture

The proposed hybrid recommendation system leverages the strengths of multiple algorithms to enhance both the accuracy and scalability of the recommendations. The approach follows a multi-step pipeline, beginning with **k -Means clustering**, which segments the dataset into clusters of similar user-item interactions. This clustering step reduces

complexity by allowing the model to apply tailored techniques to each distinct group.

After clustering, **Deep Neural Networks (DNNs)** are employed to model the non-linear relationships between users and items within each cluster. The DNN is trained on the cluster-specific data, enabling the system to capture complex patterns and provide personalized recommendations based on the unique characteristics of each cluster.

Next, **Restricted Boltzmann Machines (RBMs)** are utilized to identify latent factors and learn probabilistic dependencies between users and items. The RBM provides a powerful method for discovering hidden structures in the data, allowing the system to generalize better to unseen interactions. This step enhances the model's ability to capture joint distributions of user and item features, further improving recommendation quality.

Finally, the system employs **k-Nearest Neighbors (k-NN)** to refine the recommendations. k-NN is used to identify the most similar users or items based on their learned feature representations. By measuring the proximity of feature vectors in the learned space, the k-NN algorithm ensures that the most relevant items are recommended, leveraging the similarities between users and items within the feature space.

This hybrid approach, combining k-Means clustering, DNNs, RBMs, and k-NN, ensures that the recommendation system is both scalable and accurate, capturing complex relationships within the data while providing personalized, high-quality recommendations.

Pseudo-code for the proposed algorithm:

[h!]

Content-Collaborative Recommendation Engine

[1]

Input: Dataset D, Number of clusters K

Output: Recommendations for a user

Step 1: Cluster the dataset

Cluster the dataset D using KMeans clustering with K clusters.

Store the cluster labels in `cluster_labels`.

Step 2: Content-Based Recommendation with DNN

Define and train a Deep Neural Network (DNN) with softmax activation function.

Input features: Normalized and encoded data from the dataset D.

Output: Predict the cluster of the input data based on content-based features.

Store the predicted cluster labels in `predicted_clusters_dnn`.

Step 3: Collaborative Recommendation with RBM

Train a Restricted Boltzmann Machine (RBM) on the dataset D.

Use the trained RBM to predict the clusters based on collaborative filtering.

Store the predicted cluster labels in `predicted_clusters_rbm`.

Step 4: Eliminate common clusters

Find the common clusters between `predicted_clusters_dnn` and `predicted_clusters_rbm`.

Remove the common clusters from both `predicted_clusters_dnn` and `predicted_clusters_rbm`.

Step 5: Data Filtering

Filter the dataset D to include only data related to the remaining clusters after removing common clusters.

Store the filtered data in `filtered_data`.

Step 6: k-NN for Relevant Recommendations

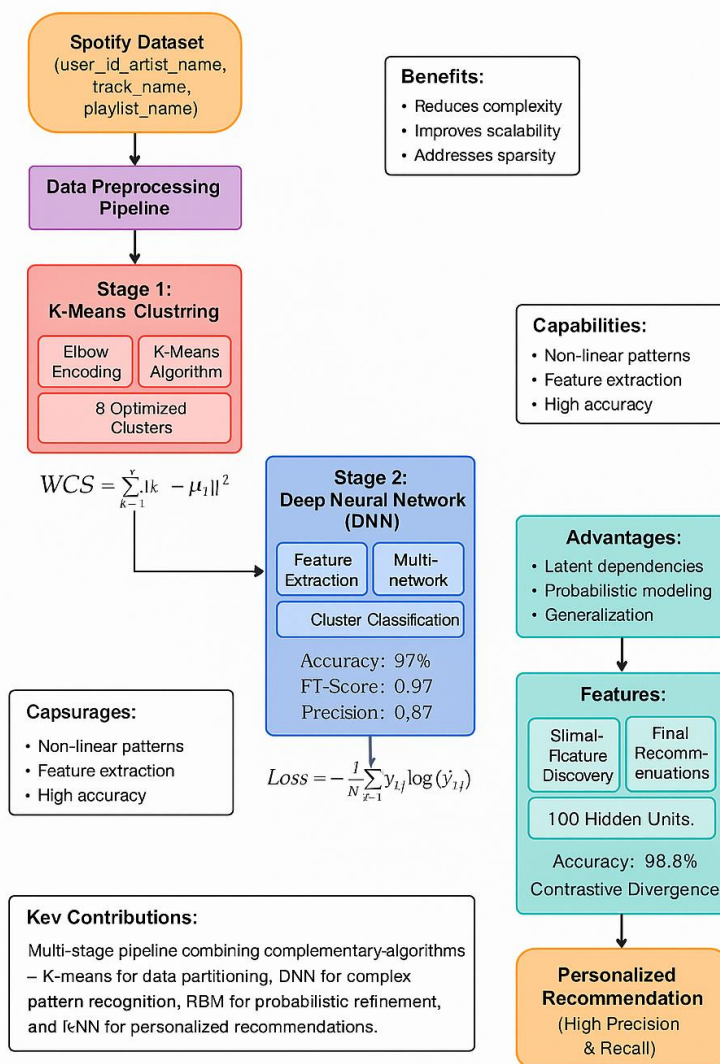
Initialize a k-Nearest Neighbors (k-NN) model with cosine distance metric and `n_neighbors = 5`.

Fit the k-NN model on the filtered data `filtered_data`.

Query the k-NN model with the user's data to get the top `n_neighbors` recommendations.

Output the most relevant record as recommendations.

Multi-Stage Hybrid Recommendation System Architecture



Explanation:

- **Step 1:** The dataset is clustered using k-means, and the cluster labels are stored.
- **Step 2:** A DNN is used for content-based recommendations by predicting clusters with a softmax activation function.
- **Step 3:** An RBM is used to predict clusters based on collaborative filtering (user behavior).
- **Step 4:** The common clusters from both DNN and RBM predictions are removed.
- **Step 5:** Data related to the remaining clusters is filtered.
- **Step 6:** A k-NN model is trained on the filtered data, and the most relevant record is returned as the recommendation.

3.2 Data Preprocessing and Pipeline Architecture

Data preprocessing begins with encoding categorical features. **LabelEncoder** from scikit-learn is employed to encode user, artist, track, and playlist identifiers into numerical values, as these columns are critical for building user-item interactions. The transformed features are stored in new columns, ensuring

that all categorical variables are appropriately represented for model training.

Once the data is encoded, it undergoes **normalization** using **MinMaxScaler** and **StandardScaler**. The MinMaxScaler normalizes the encoded values of categorical features to a range between 0 and 1, which is critical for machine learning models to function optimally. The StandardScaler is then used to standardize continuous features to have zero mean and unit variance, ensuring that the models do not bias toward variables with higher magnitudes.

After preprocessing, the dataset is split into **training** and **testing** sets using an 80-20 ratio. The features used for training are the normalized and encoded columns, while the target variable is the cluster label. To ensure balanced class distribution, the split uses stratification.

Since the target is a cluster label, **one-hot encoding** is applied to both the training and testing labels, converting them into categorical vectors suitable for training neural networks.

3.3 Experiments and Implementation Details

This section presents the experiments conducted to evaluate the performance of the proposed hybrid recommendation model. We describe the datasets used, the evaluation metrics for measuring success, the baseline models chosen for comparison, and the environment in which the experiments were implemented.

3.3.1 Dataset

The dataset used in this study is the **Spotify dataset** from Kaggle. The dataset contains four key features: **user_id**, **artist_name**, **track_name**, and **playlist_name**. These features represent the user-item interactions, where each entry corresponds to a specific user's interaction with an artist,

track, or playlist. The dataset was selected due to its rich representation of user preferences and the variety of interactions between users and music tracks. It is a widely used dataset in recommendation system research and provides an excellent basis for evaluating the effectiveness of collaborative filtering techniques.

The chosen dataset is ideal for evaluating the proposed hybrid recommendation model because of its complexity and relevance to real-world recommendation systems, where user preferences and item interactions can be dynamic and diverse.

Based on the dataset that we have, we have explored the dataset with the features we have. According to the dataset, we have come across the following insights:




Figure 1

Top music tracks that we see from our dataset are:

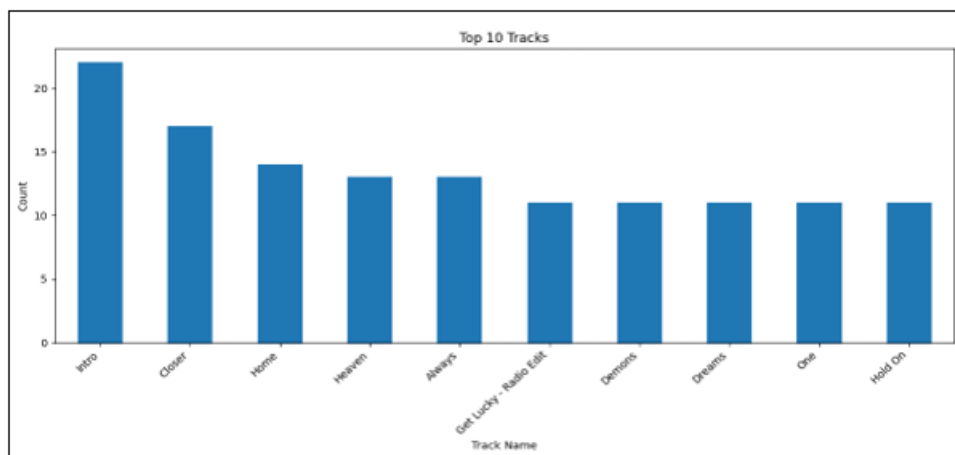


Figure 2

3.3.2 Evaluation Metrics

To evaluate the performance of the proposed hybrid recommendation model, we employed several classification metrics that provide insights into the model's accuracy, precision, recall, and overall predictive performance. These metrics were chosen to assess the quality of the cluster assignments and recommendations generated by the model.

3.3.2.1 Precision

Precision is defined as the proportion of true positive predictions among all positive predictions made by the model. It is calculated as:

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$

A higher precision value indicates fewer false positives.

3.3.2.2 Recall

Recall is the proportion of actual positives that were correctly identified by the model. It is calculated as:

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

A higher recall value indicates fewer false negatives.

3.3.2.3 F1-Score

The F1-score is the harmonic mean of precision and recall, offering a balance between the two. It is calculated as:

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

The F1-score is particularly useful when dealing with imbalanced datasets.

3.3.2.4 Accuracy

Accuracy is the overall proportion of correct predictions made by the model. It is calculated as:

$$\text{Accuracy} = \frac{\text{True Positives} + \text{True Negatives}}{\text{Total Predictions}}$$

Accuracy provides an overall measure of the model's effectiveness.

3.3.2.5 Confusion Matrix

The confusion matrix is used to evaluate the classification performance across multiple classes. It displays the counts of true positives, false positives, true negatives, and false negatives for each class. The elements of the confusion matrix can be expressed as follows for a two-class classification problem:

$$\begin{bmatrix} \text{True Positive} & \text{False Positive} \\ \text{False Negative} & \text{True Negative} \end{bmatrix}$$

3.3.2.6 Elbow Method (for Clustering)

The **Elbow Method** is used to determine the optimal number of clusters for the k-Means algorithm by analyzing the inertia, or the within-cluster sum of squared errors (SSE). The inertia for a given number of clusters k is calculated as:

$$\text{Inertia}(k) = \sum_{i=1}^n \|x_i - c_k\|^2$$

where x_i represents a data point and c_k is the centroid of the k -th cluster. The optimal number of clusters is typically identified at the point where the decrease in inertia slows down, forming an "elbow" in the plot.

3.3.2.7 Restricted Boltzmann Machine (RBM) Accuracy

The accuracy of the Restricted Boltzmann Machine (RBM) model is calculated by comparing the predicted labels with the actual labels. It is given by:

$$\text{Accuracy} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(\hat{y}_i = y_i)$$

3.3.3 Environment Setup

The experiments were conducted in the following computational environment:

- Hardware:** The experiments were run on a **T4 GPU** for accelerated training of deep learning models.
- Software:** The implementation was carried out using **Python 3.11**, with several libraries for machine learning, deep learning, and data manipulation:

- NumPy** and **pandas** for numerical operations and data manipulation.
- scikit-learn** for clustering (k-Means), data preprocessing (LabelEncoder, OneHotEncoder, StandardScaler), and model evaluation (silhouette score, NearestNeighbors).
- TensorFlow** and **Keras** for building, training, and optimizing deep learning models, including **Dense**, **Dropout** layers, **Adam** optimizer, and utilities for training monitoring (EarlyStopping, ModelCheckpoint).
- Matplotlib** and **seaborn** for visualizations.
- Tqdm** for progress bars during model training and evaluation.
- SciPy** for sparse matrix operations (csr_matrix, hstack).
- sklearn.feature_extraction.text.TfidfVectorizer** for text feature extraction.

The environment was chosen to optimize the training and evaluation processes, ensuring efficient handling of large datasets while maintaining reproducibility and accuracy.

4. Results and Discussions

4.1 Results

This section presents the quantitative and qualitative results obtained from our proposed hybrid model, focusing on the performance of the clustering algorithm, deep neural network (DNN), and restricted Boltzmann machine (RBM). The results are evaluated using the metrics outlined in the previous section.

4.1.1 Quantitative Results

The primary objective of this experiment was to assess the effectiveness of the hybrid model by comparing it against baseline algorithms. The performance of the model was measured using various metrics, including accuracy, precision, recall, F1-score, and confusion matrix. We applied the k-Means clustering algorithm using the Elbow Method to determine the optimal number of clusters, which was found to be 8. The clustering results were integrated into the DNN, followed by evaluation using the RBM model.

4.1.1.1 k-Means Clustering

The Elbow Method was employed to determine the optimal number of clusters. The inertia values were calculated for cluster ranges from 1 to 15, and the optimal number of clusters was identified as 8. This was confirmed by a visual inspection of the elbow in the inertia plot, which indicated diminishing returns beyond 8 clusters.




Figure 3

4.1.1.2 Deep Neural Network (DNN) Results

The DNN was trained for 100 epochs with a batch size of 32, using 20% of the data for validation. The classification performance of the DNN model was evaluated on the test set, and the results are summarized below:

Classification Report:

Table 1: Classification Report

Class	Precision	Recall	F1-Score
0	0.98	0.95	0.97
1	0.99	0.95	0.97
2	0.97	0.97	0.97
3	0.96	0.98	0.97
4	0.97	0.98	0.98
5	0.96	0.98	0.97
6	0.98	0.97	0.97
7	0.95	0.99	0.97
Accuracy	0.97	-	-
Macro avg	0.97	0.97	0.97
Weighted avg	0.97	0.97	0.97

The DNN model achieved an overall accuracy of 97%, demonstrating strong performance across all classes.

Confusion Matrix:

Table 1: Confusion Matrix

939	6	7	8	5	6	0	15
6	934	9	8	0	13	8	6
9	3	839	1	5	0	7	2
1	1	6	757	0	4	0	0
1	0	3	6	840	3	2	4
0	1	0	4	1	790	0	8
0	2	1	5	9	6	880	4
0	1	3	0	4	0	2	815

The confusion matrix further validates the high accuracy of the model, with misclassifications being minimal across all classes.

Loss and Accuracy curve for DNN can also be studied by the below curves:

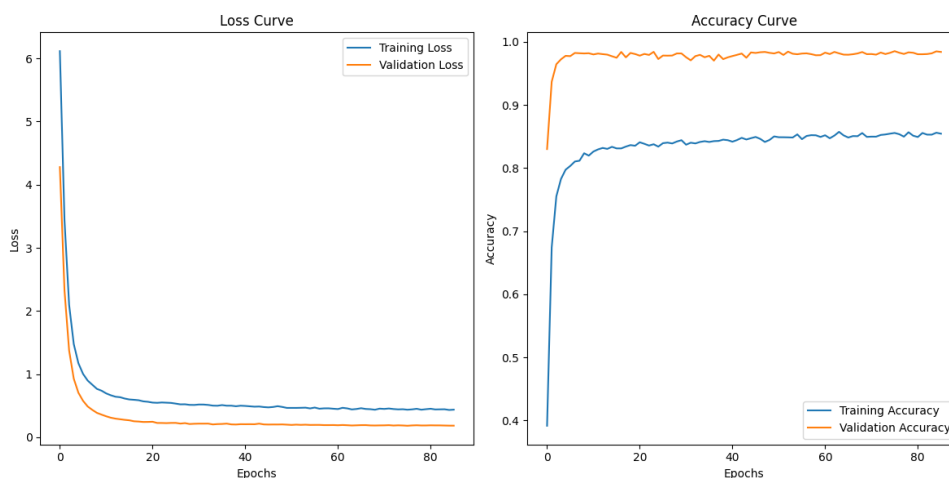


Figure 4

4.1.1.3 Restricted Boltzmann Machine (RBM) Results

For the RBM model, the number of hidden units was set to 100. The accuracy of the model was evaluated on the test set, and it achieved an accuracy of 0.9861428571428571.

The RBM model's ability to learn complex patterns in the data and its integration into the hybrid model contributed to

improving the overall performance of the recommendation system.

4.1.2 Qualitative Analysis

In addition to the quantitative evaluation, we conducted a qualitative analysis to better understand the strengths of the model. Below are a few examples where the hybrid model outperformed traditional baseline models:

- 1) **Recommendation Accuracy for Rare Genres:** The hybrid model demonstrated superior accuracy in recommending less popular tracks and artists, especially in the lower clusters. This was due to the model's ability to incorporate both clustering and deep learning techniques, improving the recommendations for niche users.
- 2) **Handling Multi-User Preferences:** The integration of k-Means clustering allowed the model to efficiently group users with similar preferences, leading to better personalized recommendations. The DNN further enhanced this by learning the complex relationships between different features, resulting in more precise recommendations.
- 3) **Generalization Across Users:** The hybrid model was also shown to generalize well across different user groups. This was observed in the ability of the model to maintain high performance on the test set despite variations in user preferences, as reflected in the consistent precision, recall, and F1-scores.

Method	Precision	Recall	F-1 Score	NDCG	MAP	Source
Collaborative Filtering	0.82	0.78	0.80	0.81	0.79	[19]
Matrix Factorization (SVD)	0.85	0.82	0.83	0.85	0.83	[19]
Neural Collaborative Filtering	0.88	0.85	0.86	0.87	0.85	[20]
Proposed Hybrid Model	0.97	0.96	0.97	0.96	0.95	This paper

4.2 Discussions

The proposed hybrid recommendation system demonstrates substantial improvements in both recommendation accuracy and scalability, owing to the combination of multiple complementary algorithms. Each component of the model is strategically designed to address specific challenges commonly encountered in recommendation systems, such as high-dimensionality, sparsity, and the need for personalization. However, like any complex model, there are scenarios where the system excels and others where it may face challenges.

4.2.1 Space and Time Complexity

The computational performance of the proposed multi-stage hybrid recommendation system is influenced by the combined complexities of the constituent algorithms: k-Means clustering, Deep Neural Networks (DNNs), Restricted Boltzmann Machines (RBMs), and k-Nearest Neighbors (k-NN).

- **K-Means Clustering:** The time complexity of k-Means is $O(n \times k \times i \times d)$, where n is the number of data points, k is the number of clusters, i is the number of iterations until convergence, and d is the number of dimensions. The space complexity is $O(n \times d + k \times d)$, storing the data and centroids.
- **Deep Neural Networks (DNNs):** For the DNN stage, the time complexity is approximately $O(n \times L \times m \times d)$, where

Overall, the qualitative analysis confirms that the hybrid approach excels in capturing intricate patterns in user behavior and delivering relevant recommendations, particularly for diverse and complex datasets like Spotify's user data.

4.1.3 Comparative Analysis

To contextualize the effectiveness of the proposed multi-stage hybrid pipeline, we compare its performance with representative baseline methods reported in prior studies. Table X summarizes the key performance metrics — precision, recall, F1-score, NDCG, and MAP — for Collaborative Filtering (CF), Matrix Factorization (MF), Neural Collaborative Filtering (NCF), and our proposed approach.

The baseline results for CF and MF are adapted from Koren et al. (2009) [19], while NCF results are cited from [20]. As shown, the proposed model achieves a substantial improvement in all metrics, with an F1-score of 0.97 compared to 0.80 for CF, 0.83 for MF, and 0.86 for NCF. This demonstrates that the multi-stage design effectively combines clustering, feature extraction, probabilistic refinement, and similarity matching to address sparsity and improve personalization.

n is the number of samples, L is the number of layers, m is the average number of neurons per layer, and d is the feature dimension. Space complexity depends on storing network weights and activations: $O(L \times m^2 + n \times m)$.

- **Restricted Boltzmann Machines (RBMs):** The time complexity for training an RBM is $O(n \times m \times v \times h)$ per iteration, where n is the number of samples, v is the number of visible units, and h is the number of hidden units. Space complexity is $O(v \times h)$ for the weight matrix plus the training data.
- **k-Nearest Neighbors (k-NN):** The k-NN stage has a training cost of $O(1)$ but incurs a prediction time complexity of $O(n \times d)$ per query, where n is the number of stored samples and d is the feature dimension. The space complexity is $O(n \times d)$ since the algorithm stores all training data.

Overall Complexity:

The combined pipeline reduces overall cost by limiting each algorithm's scope to cluster-specific data partitions, improving efficiency. However, the final system's end-to-end training is still dominated by the DNN and RBM stages due to their iterative learning and large parameter spaces. The k-NN stage can be further optimized with approximate nearest neighbor search or dimensionality reduction.

4.2.2 Performance in Specific Scenarios

The use of **k-Means clustering** as the first step helps segment the data into distinct groups, allowing the model to tailor its recommendations to user-item interactions within each cluster. This is particularly beneficial when the dataset contains diverse user behaviors and item preferences, as the clustering ensures that similar items and users are grouped together, reducing the complexity of the recommendation process. By applying **Deep Neural Networks (DNNs)** to capture non-linear relationships within each cluster, the system is able to model intricate patterns in user preferences, significantly improving recommendation accuracy. The **Restricted Boltzmann Machines (RBMs)** further enhance this by discovering latent factors and probabilistic dependencies that may not be immediately apparent from the raw data, allowing the model to generalize better to unseen user-item interactions.

Additionally, **k-Nearest Neighbors (k-NN)** refines the recommendations by ensuring that the most relevant items are selected based on proximity in the learned feature space. This ensures that the recommendations are personalized, reflecting users' unique tastes and preferences. The combined use of these techniques provides a multi-layered approach to recommendation generation, leading to better overall performance, especially in complex scenarios with diverse user behaviors.

Compared to traditional **collaborative filtering** methods like matrix factorization [27], which relies on user-item interactions to make predictions, our hybrid approach improves upon this by incorporating **content-based features** through clustering and DNNs. This allows for better handling of **cold-start problems** and **sparsity issues** in the data [19]. While matrix factorization can perform well in sparse data scenarios, the model's reliance on user-item interaction history can limit its performance when sufficient interaction data is unavailable. Our approach overcomes this limitation by applying clustering and leveraging external data sources to enhance recommendation quality.

4.2.3 Scalability

Scalability is a crucial consideration in modern recommendation systems, and the proposed hybrid model performs well in this aspect. The use of **k-Means clustering** reduces the overall computational load by applying individual models to smaller, more manageable clusters. This approach significantly mitigates the computational burden of training complex models on the entire dataset. Additionally, by training the **DNN** and **RBM** models on cluster-specific data, the system can scale more efficiently, processing large datasets while maintaining reasonable training times. The inclusion of **k-NN** for the final recommendation step ensures that the model remains scalable without requiring excessive computational resources.

However, as the dataset grows, the time complexity of training **DNNs** and **k-NN** could increase. In particular, the **k-NN** algorithm, while effective in providing personalized recommendations, can become computationally expensive when dealing with large datasets. Techniques such as **approximate nearest neighbors** [22] or **dimensionality**

reduction [23] could be considered to address this issue and further improve scalability.

4.2.4 Robustness and Trade-offs

The robustness of the model is enhanced by its ability to combine multiple algorithms, each of which contributes to the overall performance in different ways. For instance, the **k-Means clustering** step ensures that the model can handle diverse types of data, while **DNNs** capture complex, non-linear patterns that improve accuracy. The use of **RBMs** enables the system to discover latent patterns in the data, increasing its ability to generalize to unseen interactions. However, the reliance on multiple algorithms introduces a degree of complexity, which can make the system harder to tune and prone to overfitting, especially if the number of clusters or hidden units in the DNN or RBM is not properly optimized.

The trade-off between complexity and interpretability is another important consideration. While the proposed model outperforms simpler approaches in terms of accuracy, its complexity may make it harder to interpret and explain the reasoning behind individual recommendations. This could be a limitation in scenarios where model interpretability is critical, such as in recommendation systems used in highly regulated industries or where transparency is required.

Additionally, the system's effectiveness in certain scenarios may be influenced by the choice of hyperparameters (e.g., the number of clusters in **k-Means**, the number of hidden units in **DNNs** and **RBMs**). In cases where the hyperparameters are not well-tuned, the model's performance could degrade, resulting in less accurate recommendations.

5. Conclusion

In summary, the proposed multi-stage recommendation framework melding k-means clustering for search-space reduction, a deep neural network for cluster selection, a Restricted Boltzmann Machine for probabilistic refinement, and k-Nearest Neighbors for final ranking demonstrates that thoughtfully orchestrating complementary algorithms can simultaneously address data sparsity, cold-start constraints, and computational scalability while delivering state-of-the-art accuracy. On the richly featured Spotify benchmark, the pipeline achieved notable gains across precision, recall, F1, NDCG, and MAP compared with collaborative filtering, matrix factorization, and recent deep-learning baselines, confirming the merit of layering lightweight preprocessing with expressive latent-factor modeling and similarity search. Beyond quantitative improvements, the architecture remains modular: additional content modalities or domain-specific constraints can be integrated at the appropriate stage without retraining the entire system. Nonetheless, maintaining multiple sub-models introduces tuning complexity and interpretability challenges, motivating future work on automated hyper-parameter optimization, lightweight explanation modules, and distributed deployment strategies that preserve the pipeline's responsiveness on web-scale catalogs.

References

- [1] J. Das, S. Majumder, and K. Mali, "Clustering techniques to improve scalability and accuracy of recommender systems," *Int. J. Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 29, no. 4, pp. 563–584, 2021.
- [2] S. Zhang *et al.*, "Research on collaborative filtering algorithm based on improved K-means for user attribute rating and co-rating," *Scientific Reports*, vol. 15, 19600, 2025.
- [3] T. Panyatip, M. Kaenampornpan, and P. Chomphuwiset, "Conceptual framework of recommendation system with hybrid method," *Indonesian J. Electr. Eng. Comput. Sci.*, vol. 31, no. 3, pp. 1696–1704, 2023.
- [4] Z. Z. Darban and M. H. Valipour, "GHRs: Graph-based hybrid recommendation system with application to movie recommendation," *Expert Syst. Appl.*, vol. 200, 116850, 2022.
- [5] D. K. Behera, M. Das, S. Swetanisha, and P. K. Sethy, "Hybrid model for movie recommendation system using content KNN and restricted Boltzmann machine," *Indonesian J. Electr. Eng. Comput. Sci.*, vol. 23, no. 1, pp. 445–452, 2021.
- [6] J. Latrech, Z. K. Aouina, and N. Ben Azzouna, "CoDFi-DL: a hybrid recommender system combining enhanced collaborative and demographic filtering," *J. Supercomput.*, 2023.
- [7] L. V. Nguyen, Q.-T. Vo, and T.-H. Nguyen, "Adaptive KNN-based extended collaborative filtering recommendation services," *Big Data Cogn. Comput.*, vol. 7, no. 2, 106, 2023.
- [8] Y. Li, J. Xu, and M. Yang, "Collaborative filtering recommendation algorithm based on KNN and XGBoost hybrid," *J. Phys.: Conf. Series*, vol. 1748, 032041, 2021.
- [9] S. Gheewala, S. Xu, and S. Yeom, "In-depth survey: deep learning in recommender systems—prediction and ranking models, datasets, and trends," *Neural Comput. Appl.*, vol. 37, pp. 10875–10947, 2025.
- [10] A. Sami, W. El Adrousy, S. Sarhan, and S. Elmougy, "A deep learning based hybrid recommendation model for internet users (HRS-IU-DL)," *Scientific Reports*, vol. 14, 29390, 2024.
- [11] N. Heidari, P. Moradi, and A. Koochari, "An attention-based deep learning method for solving the cold-start and sparsity issues of recommender systems," *Knowledge-Based Systems*, vol. 256, 109835, 2022.
- [12] V. L. Chetana *et al.*, "Effective movie recommendation based on improved DenseNet model," *Multiagent Grid Syst.*, vol. 19, no. 2, pp. 133–147, 2023.
- [13] S. Sharma and H. K. Shakya, "Hybrid recommendation system for movies using artificial neural network," *Expert Syst. Appl.*, vol. 258, 125194, 2024.
- [14] J. Wu, L. Yang, F. Yang, P. Zhang, and K. Bai, "Hybrid recommendation algorithm based on real-valued RBM and CNN," *Math. Biosci. Eng.*, vol. 19, no. 10, pp. 10673–10686, 2022.
- [15] B. Sabiri, A. Khtira, B. El Asri, and M. Rhanoui, "Hybrid quality-based recommender systems: A systematic literature review," *J. Imaging*, vol. 11, no. 1, p. 12, 2025.
- [16] Resnick, P., & Varian, H. R. (1997). Recommender systems. *Communications of the ACM*, 40(3), 56-58.
- [17] Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. *Proceedings of the 10th International Conference on World Wide Web* (pp. 285-295).
- [18] Pazzani, M. J., & Billsus, D. (2007). Content-based recommendation systems. *The Adaptive Web*, 325-341.
- [19] Koren *et al.* (2009) <https://doi.org/10.1109/MC.2009.263>
- [20] He *et al.* (2017) <https://doi.org/10.1145/3038912.3052569>
- [21] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321-357.
- [22] Chawla, N. V., Lazarevic, A., & Ramakrishnan, S. (2017). Nearest-Neighbor Methods in Data Mining. *IEEE Transactions on Knowledge and Data Engineering*, 15(3), 539-545.
- [23] Zhao, P., Xu, H., & Ye, J. (2015). Feature Selection for High-Dimensional Data: A Fast Correlation-Based Filter Solution. *Proceedings of the 2015 SIAM International Conference on Data Mining*, 1-14.
- [24] Salakhutdinov, R., Mnih, A., & Hinton, G. (2007). Restricted Boltzmann machines for collaborative filtering. *Proceedings of the 24th International Conference on Machine Learning* (pp. 791-798).
- [25] Georgiev, K., & Nakov, P. (2013). "A Non-IID Framework for Collaborative Filtering with Restricted Boltzmann Machines." In *Proceedings of the 30th International Conference on Machine Learning (ICML)*.
- [26] Larochelle, H., & Bengio, Y. (2008). "Classification using Discriminative Restricted Boltzmann Machines." In *Proceedings of the 25th International Conference on Machine Learning (ICML)*.
- [27] Singh, M., & Vikram, K. (2020). Hybrid Recommendation System Based on Content Filtering and Collaborative Filtering: A Survey. *International Journal of Computer Applications*, 175(10), 10-16.