

# A Hybrid Music Recommendation System Based on K-Means Clustering and Multilayer Perceptron

Rafael Cintra de Araujo, Victor Moises Silveira Santos, Joao Fausto Lorenzato de Oliveira and Alexandre M. A. Maciel

*Universidade de Pernambuco, Recife - PE, Brazil*

**Keywords:** Music Recommendation Systems, Hybrid Models, K-Means Clustering, Multilayer Perceptron (MLP), Similarity Matrix.

**Abstract:** Music recommendation systems have become indispensable tools for enhancing user experiences by offering personalized playlists tailored to individual preferences. However, traditional recommendation approaches often struggle with challenges such as accurately capturing user tastes, maintaining diversity in recommendations, and addressing the cold-start problem, where limited user data hampers effective predictions. To address these issues, this study presents a hybrid recommendation model that integrates K-Means clustering and a Multilayer Perceptron (MLP) neural network to deliver coherent and diverse music recommendations. The model utilizes the all-MiniLM-L6-v2 embedding, a powerful sentence-transformer, to analyze semantic similarities in textual data such as song titles, artist names, and lyrics, encoding them into a dense vector space. Combined with normalized audio features, these embeddings enable clustering and similarity-based recommendations. Extensive experiments, conducted on datasets from Spotify and Kaggle, employed key metrics such as accuracy, F1 score, silhouette score, and cosine similarity to evaluate performance. The results highlight the system's ability to maintain genre coherence and acoustic feature consistency, minimize track repetition, and foster user engagement. Addressing challenges like the cold-start problem and diverse user preferences, the proposed model demonstrates its potential for real-world applications. Future extensions include incorporating user feedback and supporting multi-session recommendations to adapt to evolving music trends, offering a robust and innovative approach to music recommendation systems.

## 1 INTRODUCTION

Recommender systems are essential for personalizing user experiences by identifying preferences and suggesting relevant content (Sharma and Gera, 2013). In the music domain, they enhance discovery by connecting users to new songs and artists (Song et al., 2012). However, challenges such as the cold-start problem and capturing implicit user preferences persist (Roberts et al., 2014). The growth of music streaming platforms has increased data complexity, making it difficult to balance diversity, coherence, and accuracy in recommendations. Traditional collaborative and content-based filtering methods often struggle with large datasets, leading to less relevant suggestions.

To address these issues, this study proposes a hybrid recommendation model integrating clustering techniques (Hartigan and Wong, 1979), neural networks (Vogels et al., 2005), and ensemble models. By

leveraging unsupervised learning (Pola et al., 2003), the system uncovers hidden patterns in large datasets while combining content-based and collaborative filtering (Goto, nd), for improved recommendation relevance and diversity. Additionally, it explores audio features and metadata (Defferrard et al., 2017) to process complex musical data, overcoming traditional limitations like high dependence on user interaction data and poor generalization to new content.

This study advances music recommendation systems by introducing a scalable, adaptable framework capable of handling large-scale, heterogeneous data while ensuring personalization and efficiency. The paper is structured as follows: Introduction outlines the study's motivation and challenges; Related Work reviews existing methodologies; Proposed Method details the hybrid model, including K-Means and MLP; Experimental Section describes datasets, evaluation metrics, and procedures; Results analyze the system's effectiveness in balancing personalization,

diversity, and computational efficiency; and Conclusions summarize findings, discuss limitations, and suggest future improvements.

## 2 RELATED WORKS

Recent advances in recommender systems have leveraged machine learning and hybrid approaches to improve personalization and scalability. Various studies address limitations in content-based and collaborative filtering, introducing more efficient and accurate methods.

The study by (Domingues et al., 2012). explores hybrid recommendation algorithms as a solution to overcome the limitations of content-based and collaborative filtering methods. By integrating multiple data sources, the system achieved a 119% increase in absolute acceptance rate (AAR) over content-based methods and a 50% improvement over usage-based approaches, with user loyalty rates (L3R) increasing by 16%.

Godinho et al., addressed scalability and accuracy by clustering user behavior patterns, improving recommendation efficiency, though specific performance metrics were not disclosed. (Wu et al., 2024) applied deep neural networks to analyze audio features and metadata, capturing implicit user preferences and significantly improving recommendation accuracy. Their model achieved an RMSE of 0.323 in warm-start scenarios, outperforming traditional methods while maintaining precision in cold-start cases. (Godinho and Vasconcelos, nd).

Chen et al. proposed an unsupervised learning approach, using advanced clustering and deep learning to enhance scalability and efficiency for large datasets. Although specific performance metrics were not reported, the study emphasized the computational benefits of clustering in large-scale platforms.

These studies highlight the evolution of recommender systems, showcasing hybrid models, deep learning, and clustering as effective techniques for improving personalization, scalability, and accuracy in music recommendation. (Yoshii et al., 2008).

The study by Yoshii et al. integrates usage data and content features within a hybrid system, leveraging a probabilistic generative model to enhance diversity and precision in recommendations. Evaluated on a Japanese music dataset, the model achieved a 93.5% precision rate, slightly below the best collaborative method (95.2%) but introduced 90% new artists in top-ranked recommendations, effectively addressing the cold-start problem (Goto, nd)

Finally, the article by (Goto, nd). presents an effi-

cient incrementally trainable probabilistic generative model. This approach combines collaborative and content-based data to overcome cold-start issues and improve artist diversity in recommendations. The system maintained high accuracy even when introducing new users and additional data, achieving a balance between precision and variety while adapting to changing datasets (Pandya, 2024).

Table 1: Comparison of Related Works.

Author	Problem	Context	Results
Domingues et al.	Long-tail, sparse datasets	Real-time hybrid system	119% AAR, 50% usage, 16% L3R
Godinho and F Vasconcelos.	Scalability, accuracy	Large user behavior datasets	High efficiency; no metrics
Zhang et al.	Improve deep learning accuracy	Audio and metadata features	RMSE 0.323; cold-start precision
Chen et al.	Scalability, personalization	Large-scale clustering	Scalable; no specific metrics
Yoshii et al.	Diversity, cold-start problem	Hybrid collaborative model	93.5% precision; 90% diversity
Goto et al.	Cold-start, artist diversity	Dynamic real-world datasets	High accuracy, adaptable

In Table 1, the characteristics of each related work are presented in a concise and objective manner, highlighting their problems, contexts, technologies, and results. In the following section, we will delve into the proposed method of this study.

## 3 PROPOSED METHOD

This study proposes a hybrid music recommendation system that combines supervised and unsupervised machine learning techniques, clustering analysis, and recommendation algorithms to deliver a personalized and diverse user experience. The steps of the proposed pipeline are outlined, along with a detailed ex-

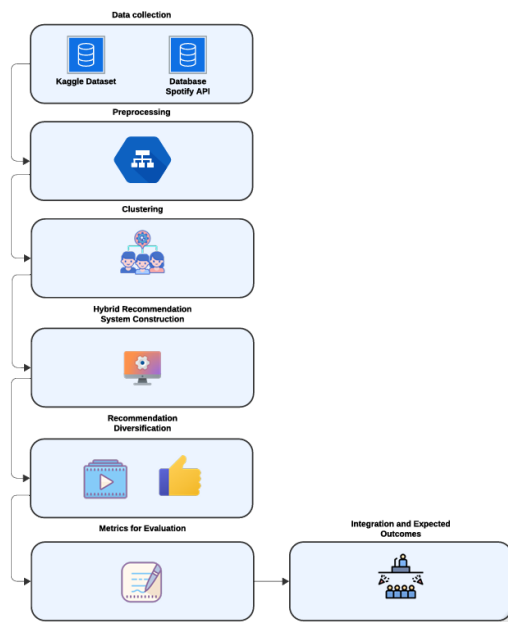


Figure 1: proposed method's flowchart.

planation of the technological decisions underpinning the approach.

Figure 1 shows the flow diagram of the proposed method, outlining the sequential steps of the system construction. Begins with data collection, leveraging both Spotify API and Kaggle datasets to enrich the available data. This is followed by preprocessing, which involves normalization and feature selection to prepare the data for analysis. The clustering phase organizes the data into meaningful groups using techniques like K-Means, forming the foundation for the next step.

### 3.1 Data Collection

The data used in this study were collected from the Spotify API, which provides detailed information on various musical attributes such as danceability, energy, valence, tempo, acousticness, liveness, loudness, instrumentalness, speechiness, mode, key, duration ms, time signature, and popularity. These attributes were selected because they comprehensively capture both the emotional and technical aspects of music, such as how users perceive energy and valence or respond to tempo and acousticness. By leveraging these diverse features, the system is designed to model user preferences more effectively and deliver personalized music recommendations. To enhance the richness of the dataset, additional information was gathered from publicly available sources on Kaggle (Pandya, 2024).

This dataset includes tracks and user interaction

data not present in the Spotify API, allowing for a broader representation of user preferences and musical diversity. The integration of external data ensures that the system captures a wider spectrum of musical genres and styles, improving the generalizability of the recommendations.

### 3.2 Preprocessing

For preprocessing (Patro and Sahu, 2015), normalization (Huang et al., 2020) was performed using the standard scaler (Aldi et al., 2023). This step ensures that all features are adjusted to a standard scale with a mean of zero and a standard deviation of one. This is particularly important for algorithms like K-Means, which rely on distance-based measures and can be affected by differences in the range of input features. Additionally, feature selection was applied to reduce the dimensionality of the data (Borges, nd), retaining only the most relevant information. This not only improves computational efficiency but also enhances the quality of the data used in the recommendation system, aligning with best practices in machine learning.

### 3.3 Clustering

The clustering process utilized the K-Means algorithm, chosen for its simplicity and efficiency in smaller-scale datasets. However, K-Means is known to face challenges when applied to large-scale data due to the computational complexity involved in distance calculations. Despite these limitations, the algorithm effectively groups songs into clusters by minimizing the variance within each cluster, creating groups of songs with similar characteristics. These clusters serve as a foundational step for generating personalized recommendations.

### 3.4 Hybrid Recommendation System Construction

The recommendation system combines both supervised and unsupervised learning methods to enhance the relevance and diversity of its recommendations. The supervised component involves the use of a Multilayer Perceptron (MLP) neural network. This model was trained to predict the cluster to which a song most likely belongs based on its features. The inclusion of the MLP was essential to address potential shortcomings of using K-Means alone on an extensive dataset, such as its limitations in handling large-scale data and its reliance on linear separability. By leveraging the MLP's ability to learn complex, nonlinear relationships and adapt dynamically to new data, the system

ensures improved accuracy and robustness in generating personalized recommendations.

### 3.5 Recommendation Diversification

Diversification is a key aspect of the proposed system, aimed at enhancing the user experience by reducing repetitive suggestions and promoting variety in the recommendations. This concept has been explored extensively in the field of recommender systems, with various state-of-the-art methods designed to balance relevance and diversity. Examples include Determinantal Point Processes (DPP), which probabilistically select diverse sets of items, and Maximal Marginal Relevance (MMR), which iteratively reduces redundancy by balancing relevance and novelty. Topic-based diversification methods are also commonly used, ensuring that recommended items span multiple categories or themes.

In the proposed system, diversification is achieved through two complementary mechanisms. The first mechanism, cluster expansion, identifies songs from neighboring clusters that are related but distinct from the user's primary profile. By incorporating these songs into the recommendations, the system maintains relevance while introducing variety. This approach encourages users to explore new music, aligning with findings in the literature that suggest unexpected but relevant recommendations can enhance user satisfaction and engagement.

The second mechanism, genre repetition penalty, addresses the issue of over representation of songs from a single genre. By applying a penalty to recommendations that disproportionately feature one genre, the system promotes a more balanced playlist. This encourages users to explore a broader range of musical styles, which has been shown to improve satisfaction and long-term retention. Together, these mechanisms create a system that balances familiarity with novelty, catering to users' evolving tastes and preferences.

### 3.6 Metrics for Evaluation

The effectiveness of the proposed system is evaluated using a combination of clustering and recommendation metrics. For clustering, the silhouette score is used to assess the quality of the clusters by measuring how well each song fits within its assigned cluster compared to others. Higher silhouette scores indicate better-defined clusters. In addition, inertia is calculated to evaluate the compactness of the clusters, with lower inertia values reflecting tighter groupings. These metrics provide a clear indication of the clus-

tering performance.

For the recommendation component, precision measures the proportion of relevant songs among the top K recommendations, while recall evaluates the system's ability to retrieve relevant songs from the total relevant set. These metrics ensure that the recommendations are both accurate and comprehensive. A diversity index is also applied to quantify the variety of the recommendations, with higher scores indicating greater diversity. For the supervised learning component, the root mean square error evaluates the prediction accuracy of the Multilayer Perceptron model, while the F1-score provides a balanced measure of precision and recall. These metrics collectively ensure a thorough evaluation of the system's performance, focusing on precision, diversity, and user satisfaction (Dehak et al., nd).

### 3.7 Integration and Expected Outcomes

The hybrid recommendation system integrates supervised learning through the Multilayer Perceptron model and unsupervised learning via K-Means clustering to achieve a balance between personalization and exploration. The supervised component facilitates recommendations that closely align with user preferences, while the unsupervised component encourages the discovery of new content through cluster analysis. By incorporating diversification mechanisms and evaluating the system with well-defined metrics, the approach demonstrates potential as a viable strategy to address the challenges of modern music recommendation systems. This combination of methods provides a foundation for creating recommendations that are both relevant and varied, addressing user needs in a flexible and scalable manner.

Figure 1 illustrates the proposed method's flowchart, outlining the sequential steps of the system's construction. It begins with data collection, leveraging both Spotify API and Kaggle datasets to enrich the available data. This is followed by pre-processing, which involves normalization and feature selection to prepare the data for analysis. The clustering phase organizes the data into meaningful groups using techniques like K-Means and PCA, forming the foundation for the next step.

The hybrid recommendation system construction integrates clustering results with supervised models to enhance recommendation accuracy and relevance. Next, recommendation diversification mechanisms ensure a balance between variety and personalization. Finally, metrics for evaluation and integration provide quantitative feedback and insights into the outcomes, guiding further refinements.

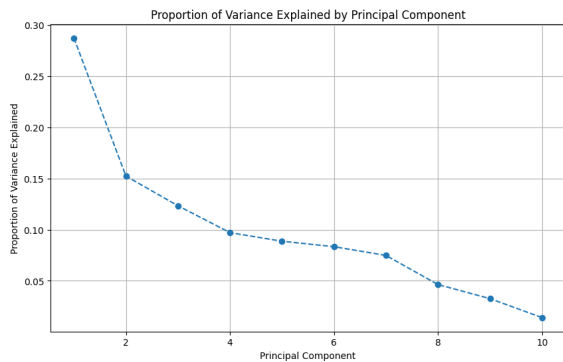


Figure 2: Elbow Method

In the following section, we will detail the experimental setup, discussing the datasets, evaluation metrics, and methodologies employed to validate the system's effectiveness.

## 4 EXPERIMENTS

The Experimental HE Section offers a thorough examination of the experimental setup and findings from evaluating the proposed hybrid music recommendation system. It covers the computational environment configuration, datasets, pre-processing steps, and employed methods. Additionally, it details the evaluation metrics, presents results in tables and graphs, and analyzes computational complexity. This section aims to highlight the ensemble approach's effectiveness and efficiency in improving recommendation accuracy.

### 4.1 Experimental Environment Configuration

The recommendation system implemented a hybrid approach, combining K-Means for initial data segmentation with a Multi-Layer Perceptron neural network to refine predictions and enhance accuracy. This integration of techniques leverages the strengths of both methods to optimize the recommendation process.

#### 4.1.1 Model and Method Settings

The K-Means algorithm was employed for initial data segmentation, partitioning data into clusters by minimizing the sum of squared distances between points and centroids. The  $n$ -clusters parameter was determined using the elbow analysis method (Cui, nd), ensuring an evidence-based selection. Hyperparameters were fine-tuned: `k-means++` improved centroid se-

lection and convergence stability, while a predefined threshold controlled termination to balance computational efficiency and clustering precision.

Figure 2 represents the Elbow Method applied to K-Means clustering, where inertia, defined as the sum of squared distances between points and their centroids, decreases as the number of clusters increases. The "elbow" point, where the rate of decline significantly slows, indicates the optimal number of clusters. This method (Cui, nd) determined that four clusters were the most suitable choice, balancing model complexity and variance explanation.

For similarity computation, the `allMiniLM-L6-v2` Transformer-based model was used, generating high-dimensional vector representations of song metadata and user preferences. Cosine similarity measured song relevance, forming the foundation for personalized recommendations.

The MLP model was structured for performance efficiency, aligning input features with segmented data to capture intra-cluster relationships. ReLU activation improved nonlinearity (Nair & Hinton, 2010), while softmax in the output layer ensured probabilistic classification. The Adam optimizer was elected for its adaptive learning rate, improving convergence stability. Categorical cross-entropy measured divergence between predicted and true distributions. These optimizations ensured a scalable and effective recommendation system.

#### 4.1.2 Metrics of Models Evaluation

Model performance was assessed using three key metrics: accuracy, categorical cross-entropy loss, and inertia.

For classification, accuracy measured the proportion of correctly classified instances, while categorical cross-entropy loss quantified the difference between predicted and actual probability distributions, guiding the Adam optimizer in refining model weights. F1 score, which combines precision (proportion of relevant recommendations) and recall (ability to retrieve all relevant songs), was particularly useful for imbalanced datasets.

For clustering, inertia measured the sum of squared distances between data points and their centroids, ensuring compact clusters. The gap statistic validated segmentation quality by comparing cluster dispersion to a random distribution. The silhouette score further assessed clustering quality, with higher values indicating well-defined, cohesive groups.

To enhance recommendation accuracy, similarity measures such as cosine similarity and Euclidean distance (Song et al., 2012) were applied to track metadata. The `allMiniLM-L6-v2` model (Bagal et al., nd)

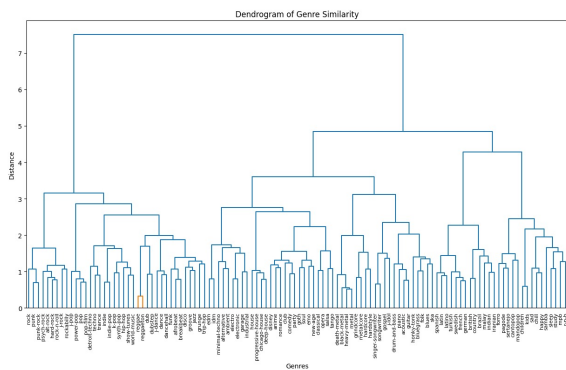


Figure 3: Dendrogram of Similarity Matrix.

transformed metadata into dense vector embeddings, forming the basis for precise recommendations. Figure 3 presents a similarity matrix (Department of Computer Science and Engineering, The Maharaja Sayajirao University of Baroda, India, 2019), visually representing genre relationships based on embedding-generated scores.

These metrics collectively provided a comprehensive evaluation framework, balancing clustering effectiveness and recommendation accuracy to refine the hybrid system.

### 4.2 Complexity Analysis

The complexity analysis of the hybrid model, integrating K-Means clustering and MLP classification, highlights its computational efficiency. While K-Means required more processing time due to its iterative centroid recalculations, it played a crucial role in structuring the dataset, streamlining the MLP classification phase.

The MLP model achieved rapid training times of 19 milliseconds per epoch, benefiting from reduced input complexity due to pre-clustered data. This efficiency resulted from its lightweight architecture and optimized configuration, enabling fast and effective learning.

Despite K-Means being computationally demanding, its contribution to organizing data into meaningful clusters significantly enhanced MLP performance. This balance between clustering complexity and classification speed showcases the hybrid system’s ability to efficiently handle large datasets while ensuring accurate music recommendations.

## 5 RESULTS

This section presents the results of the hybrid music recommendation system, evaluating its performance

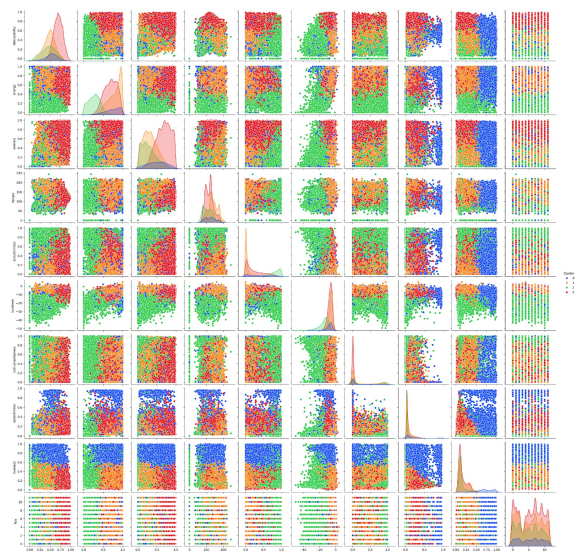


Figure 4: Pairplot of clusters.

and prediction quality. Metrics such as accuracy, clustering effectiveness, and similarity scores assess the system’s ability to provide relevant recommendations. Visualizations and comparisons illustrate the contributions of K-Means, MLP, and the similarity matrix in improving recommendation quality.

Table 2: Comparison of Features.

Feature	Average Rate of Change Between Features
Danceability	0.115
Energy	0.070
Valence	0.222
Time	28.635
Acousticness	0.052
Loudness	0.938
Instrumentalness	0.183
Speechiness	0.084
Liveness	0.140
Key	1.042
<b>Total Media</b>	<b>3.148</b>

A table summarizes the average rate of change across musical features. Acousticness (0.052) and energy (0.070) exhibit the lowest variations, indicating stable characteristics across tracks, which helps maintain smooth playlist transitions. Conversely, time (28.635) and key (1.042) show the highest variation, reflecting diversity in track duration and tonal structure. These findings highlight the system’s ability to balance cohesion and diversity.

The study developed a recommendation model that ensures personalized playlists with high coherence between genres. Cosine similarity measured the




Figure 5: Violin plot of clusters.

proximity of recommended tracks, both in sequence and in relation to the initial song. Variation analysis showed minimal differences between consecutive tracks (0.0000414) and the initial song (0.00044), ensuring smooth transitions and a consistent musical narrative.

Further analysis of danceability, energy, tempo, and valence confirmed well-controlled variability, with tempo and key exhibiting the highest fluctuations. Boxplot visualizations demonstrated an average feature variation of 3.148, validating the model’s ability to balance diversity while maintaining alignment with the initial track.

A key finding from the repetition rate analysis showed that in extensive playlists (2000 songs), previously known tracks appeared at a rate of 1.10%, reinforcing the system’s exploration mode, similar to Spotify’s recommendation feature. This ensures new recommendations while preserving structural similarity to user preferences, minimizing excessive repetition. The model effectively balances coherence and diversity, making it well-suited for music discovery applications.

Figure 4 presents a pairplot visualization, illustrating feature relationships across clusters. It highlights clear separations in features like danceability and energy, while valence and acousticness exhibit more overlap, providing insights into clustering complexity. Figure 5, a violin plot, visualizes the normalized distributions of key musical features (excluding tempo). It integrates density with a box plot summary, showing that energy and loudness have tighter distributions, while instrumentalness and acousticness display greater variation, reflecting musical diversity. These visualizations offer a deeper understanding of the dataset’s acoustic and musical properties.

## 6 CONCLUSIONS

This study introduced and evaluated a hybrid music recommendation system integrating K-Means cluster-

ing and an MLP neural network, successfully generating personalized playlists that balanced genre coherence and track diversity. Evaluations using accuracy, silhouette score, and similarity metrics confirmed the system’s robustness, enabling smooth transitions between tracks while maintaining low repetition (1.10% overlap in 2000 recommendations), making it ideal for music exploration applications like Spotify’s “Discover” feature.

A comparative analysis with Yoshii et al.’s hybrid system contextualized these results, with Table 3 highlighting key recommendation accuracy metrics. The system’s ability to seamlessly transition between genres while preserving consistency was reinforced by correlation matrices (all-MiniLM-L6-v2) and cosine similarity, validating the relevance of recommendations and demonstrating the effectiveness of hybrid models in personalization.

Table 3: Recommendation Accuracy.

Ranking (x)	Our Method (Percent)	Yoshii et al. (Percent)
Top 1	99.99%	93.5%
Top 3	99.95%	86.4%
Top 10	99.38%	80.7%

Table 4: Recommendation Diversity/Overlap.

Metric	Our Method	Yoshii et al.
Overlap with user dataset	1.10%	Not Reported
Diversity in artist/genres	High (Consistency with Genre Proximity Matrix)	High (Based on Collaborative Filtering)

Future improvements include expanding the dataset to incorporate diverse genres, languages, and user demographics for better generalization. Integrating user feedback loops (e.g., ratings, skip behavior) could improve adaptability, while transformers and graph neural networks could enhance track-user relationship modeling. Context-aware recommendations (e.g., mood, time, location) would create more dynamic playlists, and multi-session recommendations would allow the system to evolve with user preferences.

Despite its success, limitations remain. K-Means clustering exhibited longer training times, affecting scalability, which could be improved through optimized clustering techniques or alternative unsupervised learning methods. Additionally, reliance on predefined audio features may limit adaptability to emerging trends, such as non-Western music styles or genre mashups.

A key constraint is the focus on single-session recommendations, whereas real-world users interact with music over time. Temporal models would be necessary to track evolving preferences, and collaborative filtering could enhance recommendations by incorporating community-driven insights. Addressing these challenges would improve scalability, robustness, and user satisfaction.

Despite these limitations, this study represents a significant step forward in hybrid music recommendation systems, effectively balancing personalization and diversity while maintaining computational efficiency. These findings establish a foundation for future research, enabling more adaptive and enriched user experiences in music recommendation.

## ACKNOWLEDGEMENTS

This paper was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brazil (CAPES) - Finance Code 001, Fundação de Amparo a Ciência e Tecnologia do Estado de Pernambuco (FACEPE), the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) - Brazilian research agencies.

## REFERENCES

- Aldi, F., Hadi, F., Rahmi, N. A., and Defit, S. (2023). StandardScaler's potential in enhancing breast cancer accuracy using machine learning. *JAETS*, 5(1):401–413.
- Bagal, V., Aggarwal, R., Vinod, P. K., and Priyakumar, U. D. (n.d.). Liggpt: Molecular generation using a transformer-decoder model.
- Borges, H. B. (n.d.). Redução de dimensionalidade em bases de dados de expressão gênica. Technical report or unpublished work.
- Cui, M. (n.d.). Introduction to the k-means clustering algorithm based on the elbow method.
- Defferrard, M., Benzi, K., Vandergheynst, P., and Bresson, X. (2017). Fma: A dataset for music analysis. *arXiv*.
- Dehak, N., Dehak, R., Glass, J., Reynolds, D., and Kenny, P. (n.d.). Cosine similarity scoring without score normalization techniques.
- Department of Computer Science and Engineering, The Maharaja Sayajirao University of Baroda, India (2019). Research advances in computer science and engineering. *International Journal of Computer Science*. et al.
- Domingues, M. A., Trevisani, P. L., Duarte, F. M. N., and Cunha, A. M. (2012). Combining usage and content in an online music recommendation system for music in the long-tail. In *Proceedings of the 21st International Conference on World Wide Web*, pages 925–930, Lyon, France. ACM.
- Godinho, A. and Vasconcelos, F. (n.d.). Recomendação musical para grupos baseada em modelo híbrido.
- Goto, M. (n.d.). Rwc music database: Music genre database and musical instrument sound database.
- Hartigan, J. A. and Wong, M. A. (1979). Algorithm as 136: A k-means clustering algorithm. *Applied Statistics*, 28(1):100.
- Huang, L., Qin, J., Zhou, Y., Zhu, F., Liu, L., and Shao, L. (2020). Normalization techniques in training dnns: Methodology, analysis and application. *arXiv*.
- Pandya, M. (2024). Spotify tracks dataset. Accessed on December 8, 2024.
- Patro, S. G. K. and Sahu, K. K. (2015). Normalization: A preprocessing stage. *International Advanced Research Journal in Science, Engineering and Technology*, 2(3):20–22.
- Pola, G., Bujorianu, M. L., Lygeros, J., and Di Benedetto, M. D. (2003). Stochastic hybrid models: An overview. *IFAC Proceedings Volumes*, 36(6):45–50.
- Roberts, A., Brooks, R., and Shipway, P. (2014). Internal combustion engine cold-start efficiency: A review of the problem, causes and potential solutions. *Energy Conversion and Management*, 82:327–350.
- Sharma, L. and Gera, A. (2013). A survey of recommendation system: Research challenges. *International Journal of Engineering Trends and Technology*.
- Song, Y., Dixon, S., and Pearce, M. (2012). A survey of music recommendation systems and future perspectives. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 395–400.
- Vogels, T. P., Rajan, K., and Abbott, L. F. (2005). Neural network dynamics. *Annual Review of Neuroscience*, 28(1):357–376.
- Wu, L., Wang, J., Xie, Y., and He, X. (2024). A survey on large language models for recommendation. *arXiv*.
- Yoshii, K., Goto, M., Komatani, K., Ogata, T., and Okuno, H. G. (2008). An efficient hybrid music recommender system using an incrementally trainable probabilistic generative model. *IEEE Transactions on Audio, Speech, and Language Processing*, 16(2):435–447.