USING THE LISTENING HISTORY OF YOUNGSTERS TO PREDICT THE FEATURES OF THE PERFECT SONG

Author: Borislav Semerdzhiev - B.N.Semerdzhiev@student.tudelft.nl
Supervisors: Dr. Maria Soledad Pera, Robin Ungruh

Introduction

- Music plays a crucial role in children's development by helping them express their identity, teaching them to belong to a culture, and developing their cognitive well-being and inner self-worth.[1, 2]
- Current music recommendation systems are not designed to cater to the children group.
- User modelling techniques that focus on the individual user have significant potential to capture children's music preferences accurately.

Research Question

Can we use the listening history of children enriched with high-level track descriptors in order to determine how the features of a song they would frequently listen to would look like

Methodology

- Embed the song features into the 2D latent space.
- Cluster the tracks to obtain groups that contain songs with similar features[3].
- Using the user's listening history, choose a cluster that will represent his music preferences the best.
- Compute his music preferences by taking the average of the song features of the tracks he has listened to inside this chosen cluster.

Experimental setup

- LFM-2B data set containing 49,423,141 listening events from children between the age of 6-17.
- PCA to reduce the feature space from 8 features to 2.
- Silhouette analysis to find the optimal number of clusters.
- K-Means to find the clustering.
- Choose the component with the highest number of songs listened to by the user.
- Evaluate the cosine similarity score between the song that was most listened to and the calculated preferences.
- Used features: Danceability, Energy, Instrumentalness, Acousticness, Tempo, Valence, Speechless, Liveness

Results

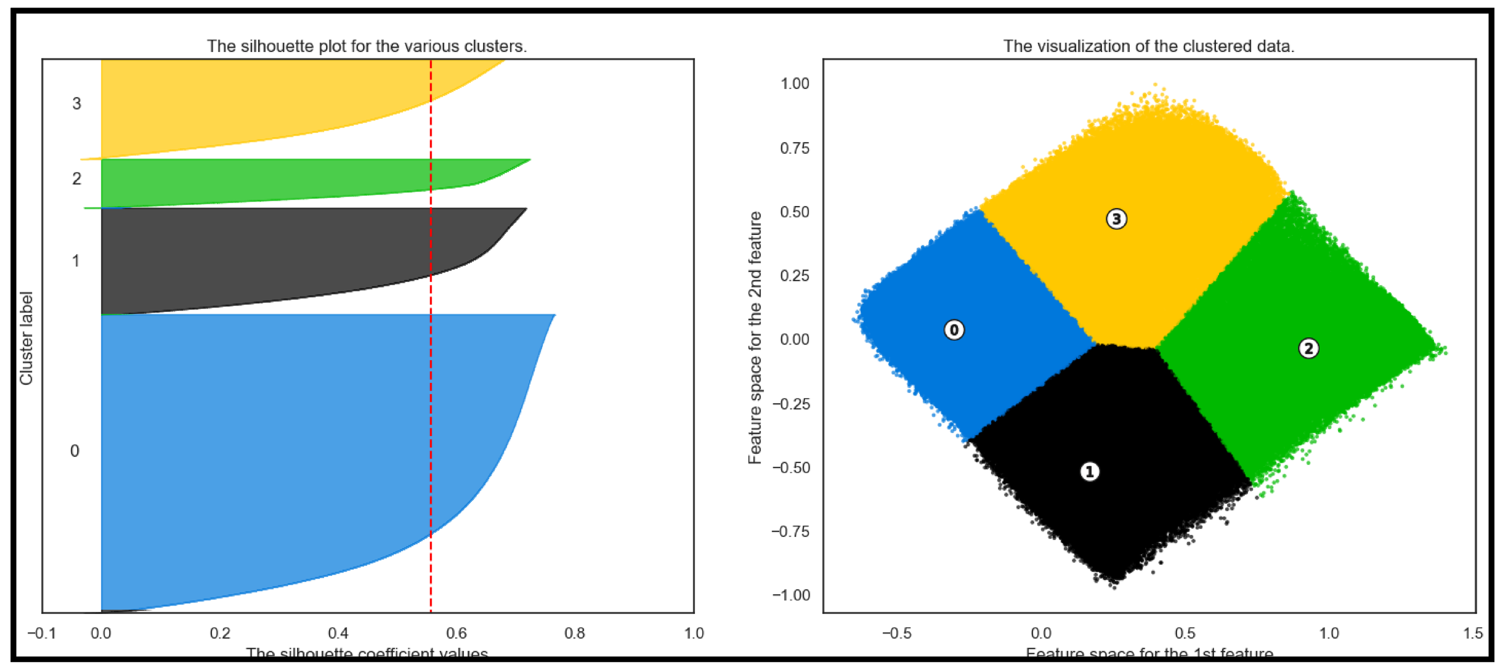

Figure 1:
Silhouette analysis of the data clustered into four components

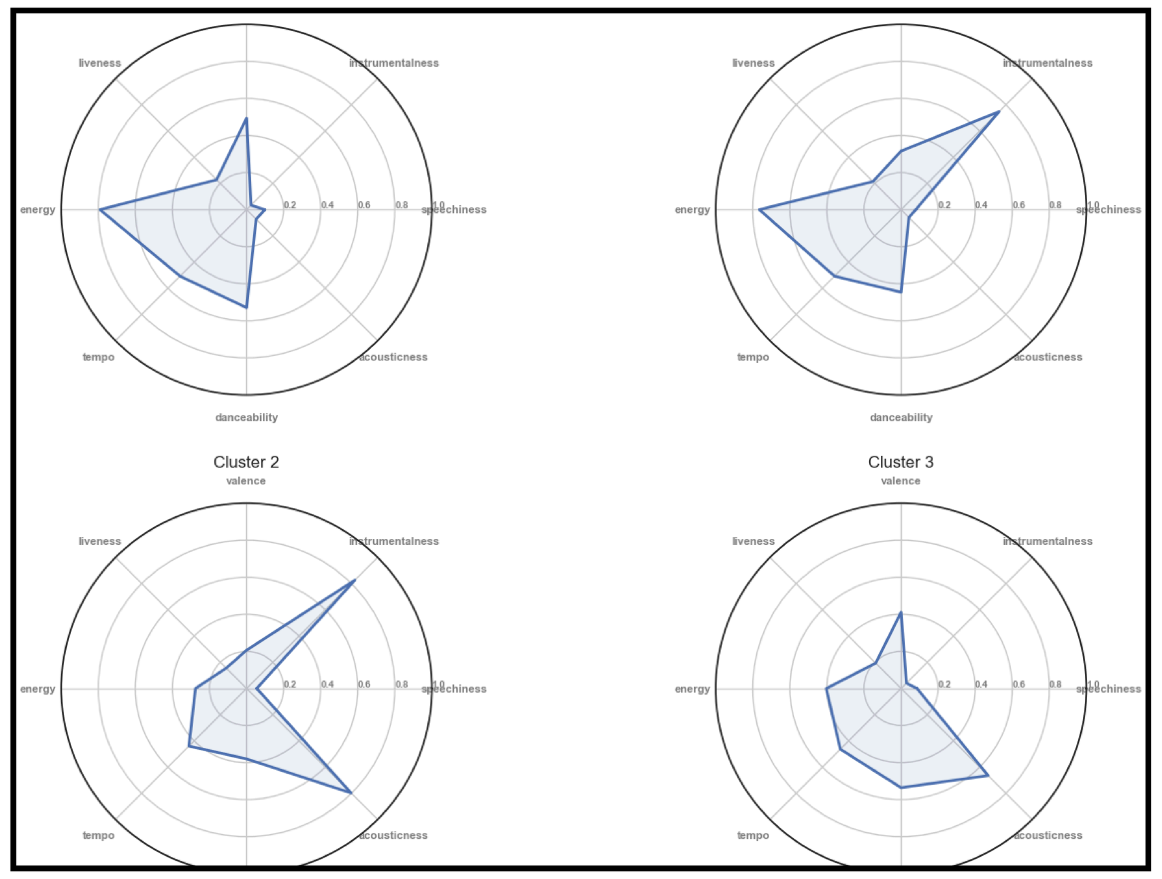

Figure 2:
Radar diagrams displaying the mean features inside each cluster

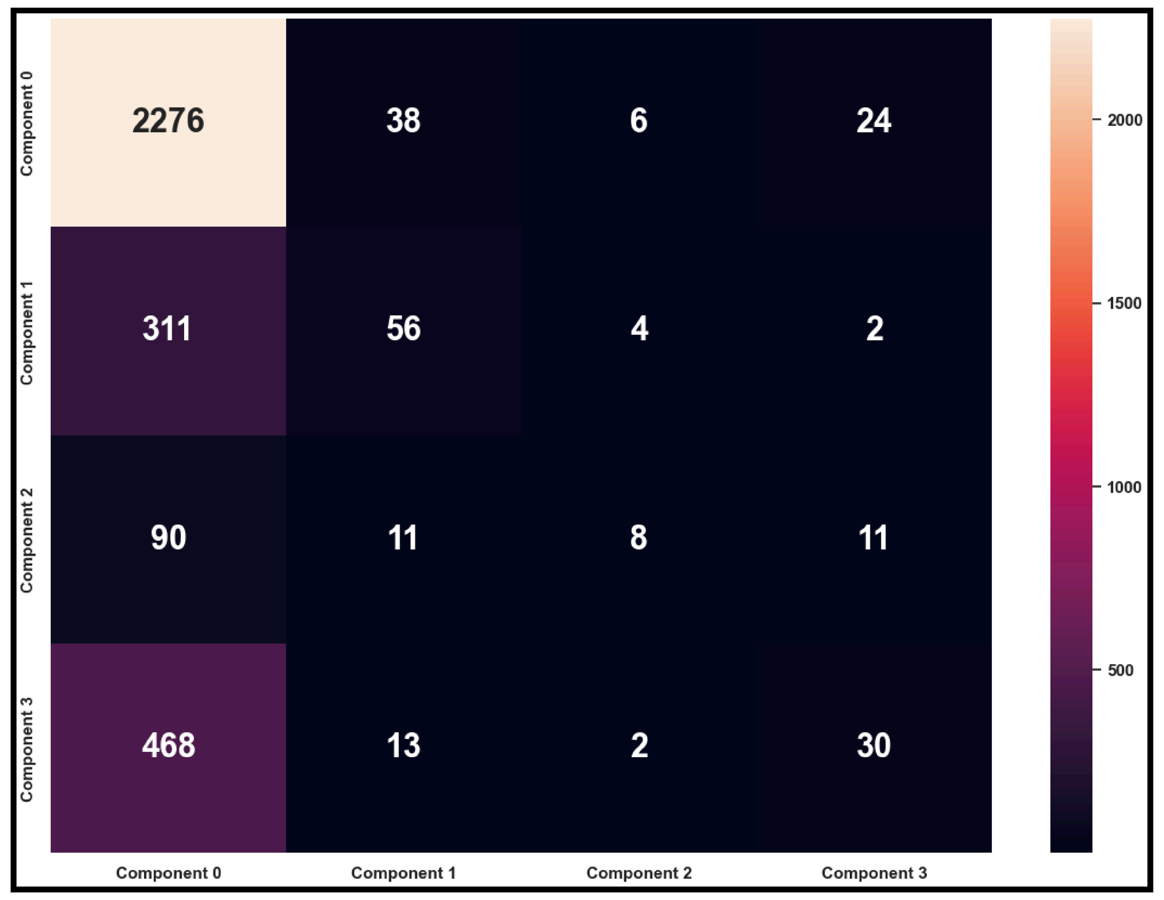

Figure 3:
Confusion matrix displaying how often we choose the wrong cluster to represent the user's preferences

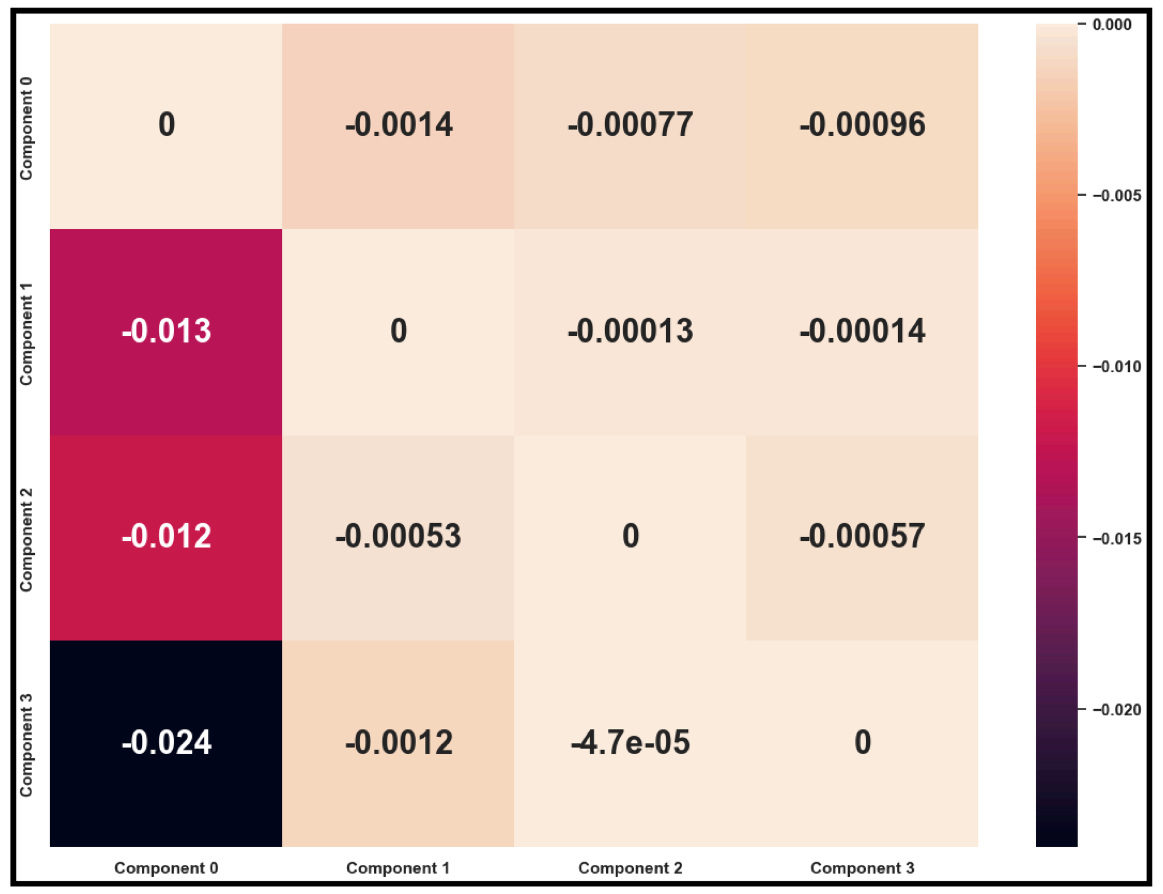

Figure 4:
Confusion matrix visualizing how much cosine similarity score we are losing due to incorrect component prediction

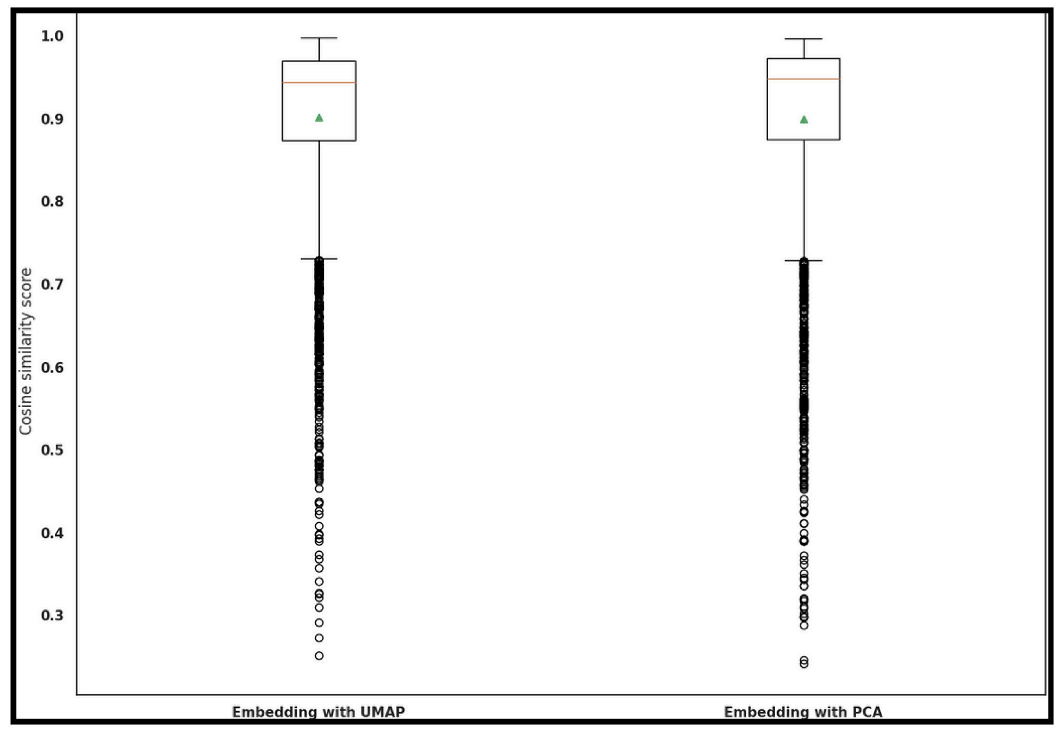

Figure 5:
Cosine similarity between the predicted preferences and the most replayed song's features

Conclusion

- Clustering of the songs based on their high-level features embedded into the 2D latent space manages to capture the similarity between different tracks.
- Selecting the component with the highest number of songs listened to by the user is an effective strategy of choosing the cluster that will represent his music preference

Findings

- Four clusters best represent the different styles of music children listen to
- There is a more popular genre of music that all the children listen to, and it contains as many songs as all the other components combined.
- Some children mainly listen to songs from the same component (Cluster 0) but have their most replayed track from another cluster.

References

[1] Susan Hallam. The power of music: Its impact on the intellectual, social and personal development of children and young people. International Journal of Music Education, 28(3):269–289, 2010
[2] Sarah J. Wilson Rudi ˇCrnˇcec and Margot Prior. The cognitive and academic benefits of music to children: Facts and fiction. Educational Psychology, 26(4):579–594, 2006
[3] Eva Zangerle and Martin Pichl. The Many Faces of Users: Modeling Musical Preference. In Proceedings of the 19th International Society for Music Information Retrieval Conference, pages 709–716. ISMIR, November 2018.