

Cluster Analysis of Musical Attributes for Top Trending Songs

Zayd Al-Beitawi
NASA JPL

zayd.b.al-beitawi@jpl.nasa.gov

Mohammad Salehan
Cal Poly Pomona

msalehan@cpp.edu

Sonya Zhang
Cal Poly Pomona

xszhang@cpp.edu

Abstract

Music streaming services like Spotify have changed the way consumers listen to music. Understanding what attributes make certain songs trendy can help services to create a better customer experience as well as more effective marketing efforts. We performed cluster analysis on Top 100 Trending Spotify Song of 2017, with ten attributes, including danceability, energy, loudness, speechiness, acousticness, instrumentalness, Liveness, valence, tempo, and duration. The results show that music structures with high danceability and low instrumentalness increase the popularity of a song and lead them to chart-topping success.

1. Introduction

Music streaming services have revolutionized the way consumers listen to music, not only by lowering the costs but also by providing consumers with an endless library of artists from all genres and musical backgrounds. As of July 2019, Spotify, the leading music streaming service, provides access to over 50 million tracks to 232 million monthly active users, including 108 million paying subscribers [16]. Spotify's payment model structures around a \$5 monthly subscription fee that provides a user with unlimited, advertising-free experience. For an additional \$5, users receive premium features including offline listening, a mobile app, enhanced sound quality, exclusive content, early album releases, and sound system compatibility [15].

In recent years, Spotify has allowed users to discover music and create exclusive playlists based on their musical preferences, favorite genres and artists, and even mood. This design has helped in eliminating a potential struggle for users in searching an extensive database of millions of songs. To optimize such discovery and personalization, streaming services like Spotify not only rely heavily on recommender systems but also on human editors [1]. A deeper understanding

of the characteristics and use of playlists and how users create and maintain their playlists can contribute to better recommendations.

As these playlists become more customized based on Spotify's recommendations, certain songs begin to recurrently appear on "Top Song" lists resulting in their trending on the platform. For each song, Spotify provides audio features such as duration, key, and mode. This study intends to investigate whether the success of the trending songs is related to these attributes. The results would allow music streaming services to create better-customized playlists that reduce search time and improve the satisfaction of their users. The findings would also lead to more focused marketing efforts by the artists to attract potential subscribers to their music.

2. Related Work

Discovery and personalization are a key part of the user experience and critical to the success of the creator and consumer ecosystem in music industry [6]. Both Content-based filtering and Collaborative filtering recommender systems were applied for discovery and personalization by both practitioners and researchers. Data scientists at Spotify had developed Discover Weekly, a personalized playlist which updates weekly and reached 1 billion streams within the first 10 weeks from its release, powered by a scalable factor analysis of Spotify's over two billion user-generated playlists matched to each user's current listening behavior [6]. Others had also generated playlist recommender systems based upon playlist names [10], social data of musicians [3], or the Facebook likes of artists and the listening history of songs of a Spotify user [4]. Finally, a survey study finds that track and artist popularity can play a dominant role in the automated playlist generation process [1]. More interestingly, a study shows that very simple popularity-based algorithms can outperform sophisticated algorithms in more general music recommendation scenarios [8].

Previous studies [5, 2, 11, 12] attempted to classify popular music data with various machine learning

algorithms, including decision tree, regression, SVM, Naïve Bayes, and neural network. Most these studies utilized a more limited and abstract set of musical attributes compared to Spotify’s audio features. Only one study [12] used Spotify’s audio features to find music popularity; the researchers conducted CART decision tree classification to a dataset containing Indonesia’s Daily TOP 200. The songs with streams more than 2 million labeled as popular and the songs with streams less than 2 million labeled as non-popular. The results found five dominated attributes represented the characteristics of popular songs - acousticness, liveness, energy, valence, and key. Songs played with acoustic instruments, medium energy, moderate valence, and high base key are considered as popular songs in Indonesia. In this study, we aim to study the similarities of trendy music in the more influential U.S. market based on Spotify’s audio features, using a different machine learning approach – clustering analysis. We hope the results from this study could contribute to discovery and personalization for consumers, as well as to music creation and promotion for creators.

3. Spotify Audio Features

Using the audio features component of the Spotify API service [14], users can extract a series of

characteristics for each song, such as how acoustic or loud it is. The list of audio features, as well as their data type and definition, are provided by Spotify as displayed in table 1.

4. Research Methodology

4.1. Dataset

At the end of each year, Spotify compiles a variety of lists showcasing the top artists, songs, and albums, and it categorizes some of the lists based on region, streaming platform, and musical genre. To analyze popular musical trends and to understand what leads to their success, we used the "Top 100 Trending Spotify Song of 2017" as our primary dataset in this study which is comprised of the top 100 most-streamed tracks on Spotify. Although we were limited to 100 records, the type of artists and genres featured on the list represent a good variability, with over five genres, as shown in Figure 1. Table 2 shows the descriptive statistics.

Then we checked out the correlations between the variables (Table 3), which are mostly consistent with their definitions in Table 1. There is a high correlation (0.71) between loudness and energy, but it will not be an issue in this study, which focuses on clustering which measures distances.

Table 1. Spotify audio features

Attribute	Data Type	Definition
Key	integer	The estimated overall key of the track. Integers map to pitches using standard Pitch Class notation. E.g. 0 = C, 1 = C#/D \flat , 2 = D, and so on.
Mode	integer	The modality (major or minor) of a track, the type of scale from which its melodic content is derived. Major is represented by 1 and minor is 0.
Time_signature	integer	An estimated overall time signature of a track. The time signature (meter) is a notational convention to specify how many beats are in each bar (or measure).
Danceability	float	Describes how suitable a track is for dancing based on a combination of musical elements, including tempo, rhythm stability, beat strength, and overall regularity. A value of 0.0 is least danceable, and 1.0 is most danceable."
Energy	float	A measure from 0.0 to 1.0 and represents a perceptual measure of intensity and activity. Typically, energetic tracks feel fast, loud, and noisy. For example, death metal has high energy, while a Bach prelude scores low on the scale. Perceptual features contributing to this attribute include dynamic range, perceived loudness, timbre, onset rate, and general entropy.
Loudness	float	An attribute of auditory sensation in terms of which sounds can be ordered on a scale extending from quiet to loud.

Attribute	Data Type	Definition
Speechiness	float	Detects the presence of spoken words in a track." If the speechiness of a song is above 0.66, it is probably made of spoken words, a score between 0.33 and 0.66 is a song that may contain both music and words (e.g. rap music), and a score below 0.33 means the song does not have any speech.
Acousticness	float	A confidence measure from 0.0 to 1.0 of whether the track is acoustic. 1.0 represents high confidence the track is acoustic.
Instrumentalness	float	Represents the number of vocals in the song. The closer it is to 1.0, the greater likelihood the song contains no vocal content.
Liveness	float	Describes the probability that the song was recorded with a live audience. A value above 0.8 provides a strong likelihood that the track is live.
Valence	float	Describes the musical positiveness conveyed by a track, with a measure from 0.0 to 1.0. Tracks with high valence sound more positive (e.g., happy, cheerful, euphoric), while tracks with low valence sound more negative (e.g., sad, depressed, angry).
Tempo	float	Describes the timing of the music or the speed at which a piece of music is played.
Duration_ms	integer	The duration of the track in milliseconds.

Table 2. Descriptive statistics

Attribute	Mean	SE	Median	SD	Kurtosis	Skewness	Range	Min.	Max.
Danceability	0.70	0.01	0.71	0.13	1.52	-0.89	0.67	0.26	0.93
Energy	0.66	0.01	0.67	0.14	-0.83	-0.33	0.59	0.35	0.93
Loudness	-5.65	0.18	-5.44	1.80	1.15	-0.88	9.07	-11.46	-2.40
Speechiness	0.10	0.01	0.06	0.10	3.51	2.00	0.41	0.02	0.43
Acousticness	0.17	0.02	0.11	0.17	1.11	1.33	0.69	0.00	0.70
Instrumentalness	0.00	0.00	0.00	0.03	45.39	6.55	0.21	0.00	0.21
Liveness	0.15	0.01	0.13	0.08	1.71	1.40	0.40	0.04	0.44
Valence	0.52	0.02	0.50	0.22	-0.66	0.04	0.88	0.09	0.97
Tempo	119.20	2.80	112.47	27.95	0.21	0.88	124.85	75.02	199.86

Table 3. Correlations

	Danceability	Energy	Loudness	Speechiness	Acousticness	Instrumentalness	Liveness	Valence	Tempo
Danceability	1.00								
Energy	-0.12	1.00							

Loudness	0.04	0.71	1.00						
Speechiness	0.09	-0.24	-0.46	1.00					
Acousticness	0.02	-0.25	-0.14	-0.05	1.00				
Instrumentalness	-0.03	0.10	-0.06	-0.09	-0.07	1.00			
Liveness	-0.07	0.13	0.05	-0.03	-0.13	-0.04	1.00		
Valence	0.38	0.31	0.42	-0.13	0.11	-0.07	-0.01	1.00	
Tempo	-0.31	0.06	-0.13	0.19	-0.24	0.15	0.06	-0.26	1.00

4.2. Cluster Analysis

Then, we conducted a cluster analysis to identify groups of trending songs with similar features. K-means clustering was used for the analysis. K-means clustering involves using “a set of n data points in real d -dimensional space, R^d , and an integer k ...to determine a set of k points in R^d ...to minimize the mean squared distance from each data point to its nearest center” [7].

Before determining the best value for k , we first cleaned our dataset and rearranged it to filter out unhelpful features. As a result, we removed track ID, song name, and artist name columns, which are all nominal and not suitable in the cluster analysis. After further visualizing the dataset, we decided also to remove the time signature column which had a low variance; it only contained time signatures of 3 and 4, which is challenging to use for more than 2 clusters. We then removed all rows with null values. Once the data was cleaned, we normalized all non-categorical values to make sure all variables have equal importance when the distance is calculated [5]. Lastly, we created dummy variables for categorical columns, which were key and mode. The genre category was not provided by Spotify and was manually collected and included in the dataset by the authors. The genre was excluded from cluster analysis and was saved for comparison with the generated clusters.

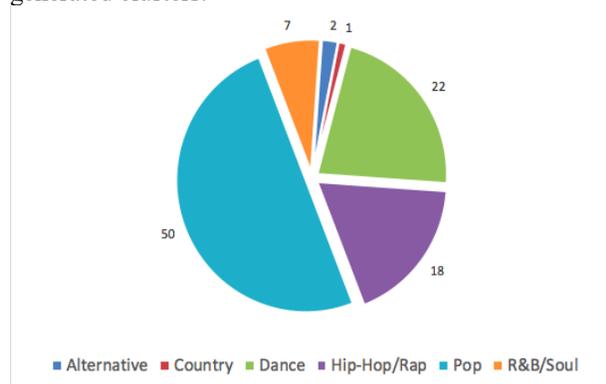


Figure 1. Spotify top 100 songs music genres

Then, we moved on to determine the optimum number of clusters for our k-means algorithm using Python programming language. We needed a set of clusters that

contained a significant amount of details without dividing up the dataset into underwhelmingly small clusters or confusingly large clusters. As a rule of thumb, we aimed at forming clusters with at least 10 records. Using the Silhouette method, 2, 4 and 5 seemed to be optimal candidates for the number of clusters as Figure 2 shows. Agglomerative clustering confirmed this view where the most gains was achieved by reducing the number of clusters to 2, 4, and 5 which increased the distance between clusters by 1.66, 1.55, and 1.41 respectively. Figure 3 shows the corresponding Dendrogram.

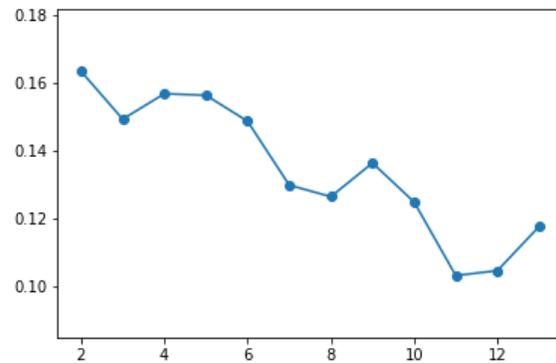


Figure 2. Silhouette chart

A cluster size of 2 was too small for analysis so it was discarded. Then, we evaluated $k=5$ which generated clusters where 2 of them has a significant overlap. In contrast, overlap between clusters was not an issue when $k=4$. Hence, we selected 4 as the optimal number of clusters and proceeded with K-means clustering. Our next objective was to characterize the clusters and analyze their patterns to determine if the top trending songs contained specific attributes that directly lead to their success. Using the established clusters, we looked at specific characteristics that result in a higher chance of trending and song types/genres that rarely make it on the top. We also wanted to see if each of the 4 clusters matched with a specific music genre, thus potentially providing us with information about the type of musical attributes that make up a specific genre.

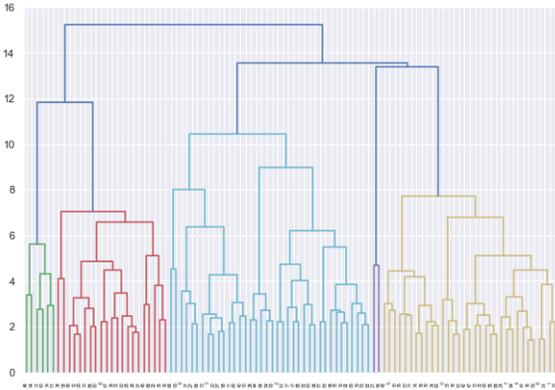


Figure 3. Dendrogram for agglomerative clustering

5. Results and Discussion

To make sense of the clusters, we drew multiple scatter plots where one dimension was a song attribute and the other dimension was the generated cluster labels (Figure 4). Among these song attributes, valence, key, and mode did not seem to significantly vary across clusters. The clustering results are summarized in Table 4. The largest cluster, Cluster#1, contained 47% of the songs and had the attributes of high danceability, high loudness, low speechiness, and low to average tempo. Songs in this cluster are upbeat, joyful, danceable, and contain fewer spoken words. Cluster#2, the second largest cluster with 27% of the songs, shares the high danceability with Cluster#1, resulting in a majority (72%) of the top trending songs having a danceable music structure. Cluster#4, which is characterized by low acousticness, average loudness, and average to high tempo, is comprised of a mix of rap, pop, and dance songs. Overall, these clusters all consist of an overwhelming majority of Pop and Dance tracks from the trending list (71 out of top 100 songs). As a result, we could conclude that the genres of Pop and Dance contain a successful, chart topping musical structure that are high in loudness and low in speechiness (Figure 5). On the other hand, the smallest cluster, Cluster #3, containing only two songs ranked at 22 and 57, presented a significant attribute - high level of instrumentalness unseen in other clusters. As the only one of the four clusters that contained high instrumentalness, this small cluster potentially emphasizes that songs with a sophisticated and varying unique musical structure, such as songs in the Alternative genre, while represent a niche market with dedicated consumers, tend to not chart as well as songs with very redundant and easy to follow beat patterns, as well as catchy hooks/phrases, as seen with the more popular trending Pop or Dance genres.

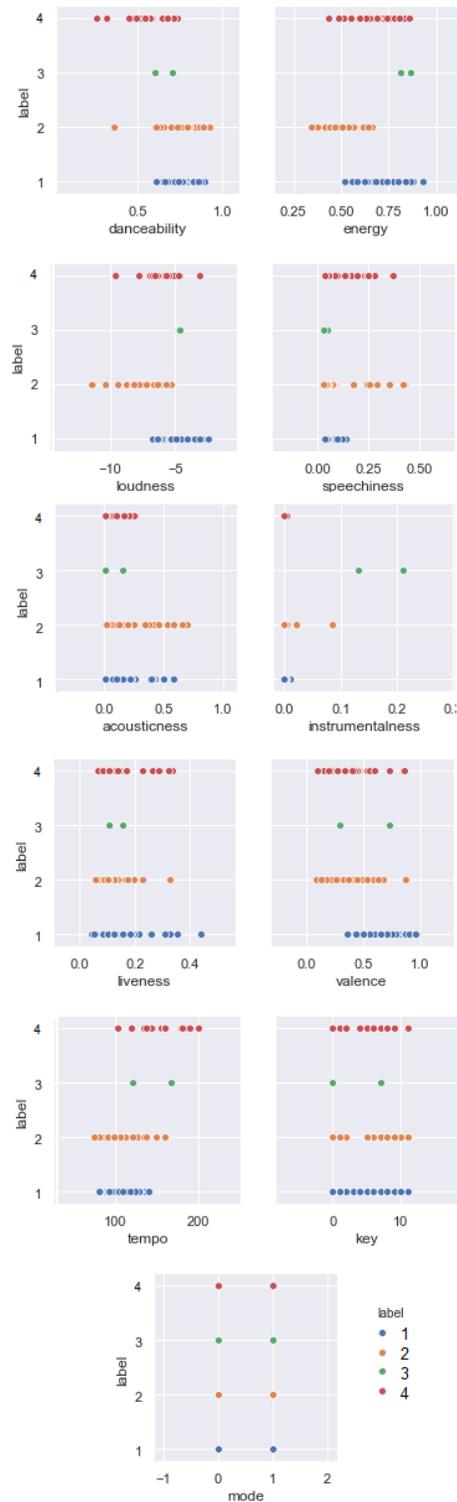


Figure 4. Distribution of clusters on different song attributes

Table 4. Cluster analysis results

Cluster #	Significant Attributes	Genres	# of Songs
1	Low Speechiness, High Danceability, Low to Average Tempo, High Loudness	Pop (26), Dance (12), R&B/Soul (2), Hip-Hop/Rap (6), Alternative (1)	47
2	High Danceability, Low to Average Energy, Low to Average Loudness, Low to Average Liveness	Pop (10), Dance (4), Hip-Hop/Rap (8), R&B/Soul (4), Country (1)	27
3	High Instrumentalness, High Energy, High Loudness, Low Speechiness, Low Acousticness, Low Liveness	Pop (1), Alternative (1)	2
4	Low Acousticness, Average Loudness, Average to High Tempo	Pop (13), Dance (6), Hip-Hop/Rap (4), R&B/Soul (1)	24

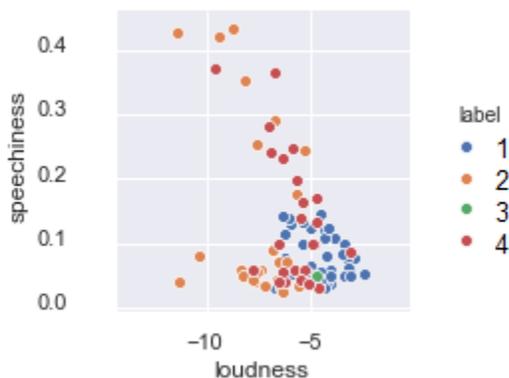


Figure 5. Speechiness vs loudness

6. Conclusion and Future Work

The intention behind conducting a cluster analysis in this study was to automatically characterize trendy music based on the musical attributes defined by Spotify. We found clusters that not only vary in size but also contain a variety of significant attributes in each cluster. The completeness and homogeneity scores between clusters and genres were equal to 7.18% and 8.26% respectively. These low scores indicate little overlap between genres and our clusters. This approach challenges the traditional music genres and provides new insight into how music can be automatically classified into different trending categories based on musical attributes and potentially provide better recommendations. The most popular songs tended to be the more exciting, radio-friendly songs that we all hear on our commute to work or while shopping at a supermarket. These songs follow a formulaic, pop-friendly sound, with a danceable music structure that tends to put audience in a good mood. Meanwhile, songs with high instrumentalness would not top the charts because although they may appeal more to people with exclusive or alternative tastes, they do not tend to attract or retain the mainstream listeners.

As future work, we will try to optimize our model and results with larger sample size, perform time series analysis and forecasting, also explore additional attribute of trendy music across genre, culture, time, and whether those vary across different segments (e.g., age, location, social-economical class, etc.). In the long run, we will create a recommender agent that provides better discovery and personalization for both consumers and creators based on musical attributes and the clusters automatically generated from popular songs in the past.

7. References

- [1] G. Bonnin and D. Jannach, "Automated Generation of Music Playlists: Survey and Experiments", ACM Computing Surveys (CSUR) Surveys, Volume 47 Issue 2, January 2015.
- [2] C. Chiru and O. Popescu, "Automatically Determining the Popularity of a Song" in Lecture Notes in Artificial Intelligence and Bioinformatics, 2017, pp. 488-491.
- [3] B. Fields, "Contextualize Your Listening: The Playlist as Recommendation Engine", A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy of the University of London. Available at: https://www.researchgate.net/profile/Benjamin_Fields2/publication/266271406_Contextualize_Your_Listening_The_Playlist_as_Recommendation_Engine/links/561e2d5108aef097132b32bc.pdf
- [4] Germain, A. and Chakareski, J., "Spotify Me: Facebook-assisted automatic playlist generation", Proceedings of the 2013 IEEE 15th International Workshop on Multimedia Signal Processing (MMSp), Pula, Italy, September 30 - October 2, 2013, pp. 25-28.
- [5] D. Herremans, D. Martens, K. Sorensen, "Dance Hit Song Prediction", Journal of Music Research, Volume 43 Issue 3, July 2014, pp. 291-302.
- [6] K. Jacobson, V. Murali, E. Newett, B. Whitman, and R. Yon, "Music Personalization at Spotify", RecSys '16 Proceedings of the 10th ACM Conference on Recommender Systems, Boston, Massachusetts, USA, September 15 - 19, 2016, pp. 373-373.

- [7] T. Kanungo, "An efficient k-means clustering algorithm: Analysis and implementation". *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 2002(7), pp. 881-892.
- [8] B. McFee, T. Bertin-Mahieux, D. Ellis, and G. Lanckriet (2012), "The Million Song Dataset Challenge". In *Proceedings of AdMIRe'12 Workshop*, Lyon, France, April 16–20, 2012, pp. 909-916.
- [9] M. Pichl, E. Zangerle, and G. Specht, (2016). "Understanding playlist creation on music streaming platforms". *Proceedings of the 2016 IEEE International Symposium on Multimedia (ISM)*, San Jose, California, December 11-13, 2016, pp. 475-480.
- [10] M. Pichl, E. Zangerle, and G. Specht, "Towards a Context-Aware Music Recommendation Approach: What is Hidden in the Playlist Name?", *IEEE International Conference on Data Mining Workshop (ICDMW)*, Atlantic City, NJ, USA, November 11-17, 2015.
- [11] R. Rajyashree, A. Anand, S. Yash, and H. Mahajan (2018), "Predicting Hit Music using MIDI features and Machine Learning", *The 3rd International Conference on Communication and Electronics Systems (ICCES)*.
- [12] L. Ramadhina Febirautami, I. Surjandari, and E. Laoh, "Determining Characteristics of Popular Local Songs in Indonesia's Music Market", *Proceedings of the 2018 5th International Conference on Information Science and Control Engineering (ICISCE)*, Zhengzhou, China, 20-22 July, 2018,
- [13] G. Shmueli, *Data mining for business analytics: concepts, techniques, and applications in R*. 2017: John Wiley & Sons.
- [14] Spotify. (n.d.). "Get Audio Features for a Track", Available at: <https://developer.spotify.com/documentation/web-api/reference/tracks/get-audio-features/>
- [15] E.J. Teague, "Saving the Spotify revolution: Recalibrating the power imbalance in digital copyright", *Case W. Res. JL Tech. & Internet*, 2012. 4: pp. 207.
- [16] Wikipedia.org. (2019). Spotify. [online] Available at: <https://en.wikipedia.org/wiki/Spotify>