

# Exploring Diverse Sounds: Identifying Outliers in a Music Corpus

Le Cai<sup>1</sup>, Sam Ferguson<sup>1</sup>, Gengfa Fang<sup>1</sup>, and Hani Alshamrani<sup>1</sup> \*

Creativity and Cognition Studios  
Faculty of Engineering and IT  
University of Technology Sydney

**Abstract.** Existing research on music recommendation systems primarily focuses on recommending similar music, thereby often neglecting diverse and distinctive musical recordings. Musical outliers can provide valuable insights due to the inherent diversity of music itself. In this paper, we explore music outliers, investigating their potential usefulness for music discovery and recommendation systems. We argue that not all outliers should be treated as irrelevant data, as they can offer unique perspectives to contribute to a richer musical understanding. We attempt to identify 'Genuine' music outliers, which may reveal unique aspects of an artist's repertoire and serve to enhance music exploration and discovery.

**Keywords:** Music Outlier · Music Outlier Detection · Audio characteristics · Music discovery

## 1 Introduction

In the field of music information retrieval, a primary focus is often on finding similarities among digital musical recordings, to enable recommendation systems and facilitate music discovery [10, 14, 23, 5]. Given this context, analysis of outliers has attracted less research attention [18], as they are often considered irrelevant data and removed during preprocessing, or are naturally scored lower by most similarity-focused algorithms [7]. However, outliers in the context of music can provide interesting insights and reveal unique patterns, as music inherently exhibits great diversity [20, 13].

In this paper, we explore the identification and categorization of music outliers, with an aim to ultimately enhance music discovery and recommendation systems. We propose a method to describe and discover genuine musical outliers based on audio characteristics, such as tempo and loudness. By doing so, we aim to identify outliers that can provide valuable information for music discovery while not being non-musical.

---

\* We would like to acknowledge the support received from Assoc. Prof. Sam Ferguson. In Addition, Le Cai wants to thank his partner Hanyu Meng for her unwavering patience and constant presence in his life.



This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

We present a definition of what constitutes a 'Genuine' music outlier and investigate its characteristics. Genuine outliers exhibit unique characteristics that set them apart from an artist's main style, providing insightful information for music discovery.

This paper is structured as follows: The introduction presents the motivation of our study, followed by a comprehensive literature review, which discusses the relevant background and prior research. Next, the aims and objectives are followed by the methodology section which outlines our proposed definition of Genuine music outliers and the subsequent dataset and algorithm developed for the detection. The results and discussion section evaluates the effectiveness of our algorithm based on the dataset and provides insights into its performance. Finally, the conclusion summarizes our findings and highlights the implications of our work, while also suggesting potential avenues for future research in the realm of music outlier detection and analysis.

## **2 Related Work**

### **2.1 What Makes a Song Different from Audio?**

Understanding outliers in the context of music recommendation systems necessitates a thorough examination of their diverse nature and the ability to differentiate actual music from other forms of audio. Müller [17] provides a comprehensive overview of music structure analysis, focusing on techniques for segmenting and organizing music into meaningful sections, laying the foundation for understanding the key aspects of music structure and demonstrating how various representations and algorithms can be used to analyze and compare music pieces. A system for finding structural descriptions of musical pieces defines the structure of a piece as segments with specific time ranges and labels, with segments sharing the same label considered occurrences of a particular structural part [21]. In another study, a multi-task deep learning framework is introduced for directly modeling structural semantic labels in music, such as "verses" and "choruses", from audio signals. This approach proposes a 7-class taxonomy that includes intro, verse, chorus, bridge, outro, instrumental, and silence, and consolidates annotations from four different datasets [27]. A large-scale analysis of songs in 315 different societies has found that songs share universal features like tonality, rhythm, and repetitive structures [16]. A study conducted by Shuqi et al [8]. analyzes the significance of repetition and structure in music, specifically in popular music, and demonstrated that deep learning models often struggle to identify these essential elements, which are crucial for generating coherent and appealing musical pieces. Subsequently, Sargent et al [22]. introduced a fourth principle: regularity. This principle posits that musical segments possess a certain degree of regularity, which can be leveraged to better understand and analyze the structure of the music.

### **2.2 Outlier Detection Approaches**

The purpose of outlier detection algorithms is to identify patterns and samples that deviate significantly from the normal characteristics of a group of data [12]. The reason to detect outliers is it can providing interesting insights contribute to a richer understanding of an artist's work. General outlier detection methods can be categorized into 4 categories based on an overview conducted by [26]:

**Clustering-based methods:** use a clustering algorithm to classify the majority of the elements of the set, while also clearly defining the outlying elements of the set.

**Density-based methods:** identifying outliers as points in low-density regions within a data set.

**Distance-based methods:** determining the distance between points, and considering outliers to be points with a large distance from their nearest neighbors.

**Statistical methods:** using measures such as mean, median and standard deviation to identify data points that fall outside of a defined range.

Clustering-based methods, exemplified by [2], utilize subspace cluster analysis to construct classification trees while addressing dataset scarcity. A deep learning model using a Clustering Augmented Learning Method (CALM) classifier improves genre classification by extracting deep time series features [11]. In contrast, density-based methods like DBSCAN [9] and OPTICS [1] identify outliers in low-density regions within a dataset, as demonstrated by the OPTICS algorithm applied to traditional Chinese folk music [28]. In the exploration of automatic outlier detection methods on music genre datasets, Lu et al [15] characterized outliers using their musical attributes, demonstrating the potential of these techniques to unveil unique insights into music structure and diversity within genre classification. Distance-based methods, including K-means [4] and CLARANS [19], examine the distance between data points and their nearest neighbors, considering outliers as points with large distances from the nearest neighbors. This approach has been successfully applied to traditional Irish music in [24]. Meanwhile, statistical-based methods employ measures such as mean, median, and standard deviation to identify data points outside a defined range to reveal patterns in cluster structure dynamics in popular music data [25]. However, while these studies focus on detecting outliers based on their statistical properties, the potential of the outliers themselves for music discovery has not been as extensively investigated.

### 3 Aims & Objectives

In this paper, our aims center around exploring the potential of music outliers for music discovery. To achieve this, we propose the following objectives:

**Propose an approach to describe musical outliers:** by examining various attributes that distinguish them from an artist's typical style, facilitating their identification and analysis. This implies establishing a clear definition to describe 'Genuine' music outliers as a distinct category of outliers that exhibit meaningful deviations from a set of existing digital musical recordings.

**Categorise music outliers:** into meaningful categories based on their distinguishing characteristics to create a meaningful interpretation, that also can help find the outliers that are helpful to understanding and discovering interesting music while excluding outliers that hold little data.

### 4 Method

We introduce the concept of genuine outliers within the context of music data and explore their potential value in recommendation systems. To accomplish this, we first

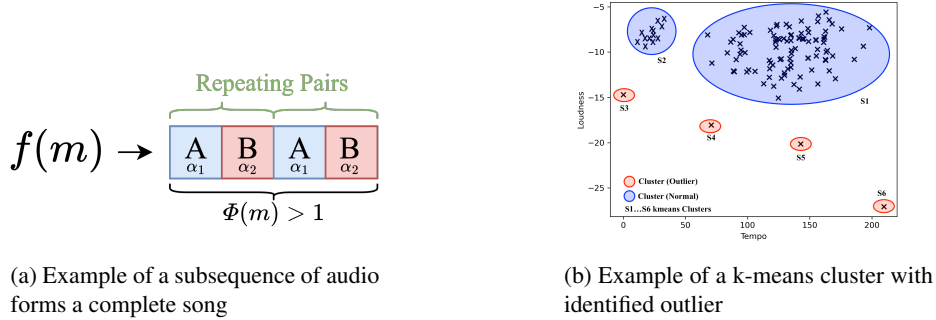


Fig. 1: The Definition of Genuine Music Outliers

propose a definition for genuine outliers, then create a labelled dataset for evaluation, and finally, apply an outlier detection algorithm to validate our definition.

#### 4.1 Definition of Genuine Music Outliers

A ‘‘Genuine’’ music outlier is a complete song that maintains an artist’s typical musical structure while distinctly diverging in sound and style from their predominant body of work, due to the complex nature of music, in our scenario, we consider pop music only.

To achieve this, a genuine outlier must satisfy the following constraints: 1): Forms a complete song, this distinguishes the identified outlier must be a song, not something else, e.g. speech, or sound effect, shown in fig.1a. To ensure this, audio must satisfy these conditions: a): The length of music structure  $\phi$  must be greater than one in a subsequence, this means the input recordings must at least have one or more music structures otherwise it is not music. b): The identified music structure must at least has one unique pair. e.g. a music structure in  $(A - B)$  is a song that can identify at least one unique pair, not  $(A - A)$  or  $(B - B)$ . c): The identified music structure must contain repeated parts, as repetition is a crucial element in music. A piece of audio that forms music should exhibit repeating sections, such as  $(A - B - A - B)$  patterns. The formal definition of a Genuine outlier is as follows:

Let  $M$  represents the set of all recordings produced by a specific artist. For a finite set  $S_M$ , suppose function  $\Phi : M \rightarrow \mathbb{N}^+$  that maps a recording to a finite integer and a function  $f_\Phi : M \rightarrow S_M^{\Phi(M)}$  that maps a recording  $m \in M$  to a finite length sequence. A recording  $m \in M$  is defined as a ‘‘Genuine’’ music outlier if the following conditions is satisfied:

**1. Forms A Complete Song:** For any recording  $m \in M$ , the sequence  $f_\Phi(m)$  and number  $\Phi(m)$  should satisfy the following:

1.  $\Phi(m) > 1$ ,
2. There exist a subsequence  $a_{k_1} \dots a_{k_p} (k_1 < \dots < k_p, p > 1)$  from  $f_\Phi(m)$  that, at least one repeated subsequence from  $f(m)$  can be found, i.e., there exist  $a_{k'_1} \dots a_{k'_p} (k'_1 < \dots < k'_p, p > 1)$  that  $k_1 \neq k'_1$  and  $k_p \neq k'_p$ .




Fig. 2: Dataset Creation Process

For any given recording  $m \in M$ , if the previous condition satisfied, for given constant  $C_G \in (0, 1]$ ,  $\kappa$  and a positive integer  $N_d$ , either of the following conditions must be satisfied,

**2. Distinctiveness:** Define  $F : M \rightarrow \mathbb{R}^n$  be a function that maps each recording to an  $n$ -dimensional feature space, where  $n$  is a positive integer representing the number of musical features being considered. Define  $\text{Card}(\cdot)$  as the Cardinality of a set,  $\|\cdot\|$  as the 2-norm and  $\kappa$ -means cluster sets  $\Omega = \{\Omega_1, \dots, \Omega_\kappa\} (1 \leq \kappa \leq \text{Card}(M))$  that

$$\begin{aligned} \Omega \in \arg \min_{\Omega} \sum_{i=1}^{\kappa} \sum_{\mathbf{x} \in \Omega_i \subset M} \|F(\mathbf{x}) - \mu_i\|^2 \\ \text{s.t. } \mu_i = \frac{1}{\text{Card}(\Omega_i)} \sum_{\mathbf{y} \in \Omega_i \subset M} F(\mathbf{y}), \text{ for } 1 \leq i \leq \kappa. \end{aligned} \quad (1)$$

For any given integer  $N_d > 0$ , for any  $m \in M$ , if  $m \in \Omega_i$  and  $\text{Card}(\Omega_i) < N_d$ , we say that recording  $m$  is distinct with respect to  $M$ .

**3. Non-adherence:** For a given positive number  $C_G \in (0, 1]$ . We define

$$R_M(m) := \frac{\text{Card}(\{m' | f_{\Phi}(m') = f_{\Phi}(m) \text{ for all } m' \in M\})}{\text{Card}(M)}. \quad (2)$$

If  $R_M(m) < C_G$ , we say  $m$  is not adhere to set  $M$ .

In conclusion, if recording  $m$  does form a complete song, either the song  $m$  is distinct with respect to  $M$  or  $m$  is not adhere to set  $M$ , with given  $C_G$ ,  $\kappa$ , and  $N_d$ , then we say song  $m$  is an outlier.

#### 4.2 Dataset Creation and Outlier Selection Algorithm

In order to evaluate our above-proposed definition for Genuine music outliers, we created a dataset specifically designed to examine our hypothesis, To achieve this goal, the following steps were carried out:

**Artist Selection:** From the Million Song Dataset <sup>1</sup>, a smaller sample consisting of 10,000 songs extracted from the MSD, we randomly chose 20 artists using a random

<sup>1</sup> Million Song Dataset [6] is available at: <http://millionsongdataset.com>

number generator to pick artist IDs. For each of the selected artists, we identified a list of all the songs contained in the dataset with the particular artist id.

**Identifying Outlier Songs:** To identify and categorize outliers within the selected artists' music, we conducted a manual listening and labeling process for these songs by ear. First, we acquainted the artist's typical styles. To achieve that, we listened to about 20%-30% songs in this artist's main cluster. After establishing the primary styles, we listened to these songs once with selected artists again, focusing on identifying tracks that significantly deviated from the typical style. Attention was paid to musical elements, such as tempo, melody, harmony, instrumentation, and song structure.

**Categorizing Outliers:** We then listened to the characteristics of the identified outliers, focusing on specific attributes such as tempo, melody, harmony, instrumentation, and song structure that differentiate them from the artist's dominant style. By examining these properties, we classified outliers into five categories: Error, Speech, Intro, Sound Effect, and Genuine. Each identified outlier was assigned to one of the categories based on a two-step process. First, we selected songs with a Euclidean distance greater than 3 times the z-score threshold from the main cluster. Second, we manually reviewed these outliers, focusing on listening to their distinctive features such as tempo, loudness, melody, and the composition of instruments.

**Creating the Labeled Dataset:** Finally, After identifying and categorizing outliers, we compiled a list of all songs from the selected artists, along with their corresponding outlier categories. This list included each song's title, artist, genre, and other metadata. Utilizing the pre-extracted audio features from the MSD Subset, such as tempo, loudness, key, and mode, we created a dataset compiled with data including outlier categories and the MSD pre-extracted features.

We consider this to be a classification problem where each artist typically consists of 1 main distinct style. Note we only considered the case for 1. Distinctiveness and 2. Non-adherences for the current approach, the case of forms a complete song is discarded due to the complexity of the analysis of music structures.

Let  $M$  be a set of recordings under an artist. Let  $x_i \in \mathbb{R}^d$  to denote a  $d$ -dimensional feature vector for the  $i$ -th song from  $M$  ( $1 \leq i \leq \text{Card}(M)$ ). Suppose  $K$  is an integer that  $1 \leq K \leq \text{Card}(M)$ . For  $1 \leq k, k' \leq K$ , define  $C_k$  as a subset of  $M$ , such that  $C_k \cap C_{k'} = \emptyset$  for  $k \neq k'$ , and  $\cup_{1 \leq i \leq K} C_i = M$ . In our case,  $d = 2$ . We adopt k-means algorithm [3] to partition set  $M$  that satisfies such conditions. In the next section, we first delineate the outlier categorization, followed by the presentation of outlier detection results.

## 5 Results

### 5.1 Outlier Categorisations

We found 34 genuine outliers from 29 artists in 320 songs as well as 17 errors, 5 speeches, 1 intro, 5 noise, and 3 sound effects from 29 artists shown in Fig. 3, they can categorize into 5 types: *a) Error:* such as erroneous data value, e.g. doubling tempo due to feature extractor failure, incorrect data type, such as text entered as a numeric type. *b) Noise:* unintentional sounds such as non-musical noise, jitter, or glitches in a




Fig. 3: Outlier Dataset: Categorisation of Outliers by Artists

recording, often due to recording issues or equipment quality. e.g. a live performance recording with excessive audience noise. c) **Speech**: such as some spoken words, or a short story integrated into the music to enhance the atmosphere of the listening experience. these types of practices commonly occur in certain music genres, such as hip and electronic music. d) **Sound Effect**: This type of recording is used to create a certain mood, it can build tension and add extra impact to certain sections of the song. Particularly, sound effect is used in certain genres of music such as electronic & rock music to add extra depth and interest to music. e) **Intro**: A very short track is often less than 30 seconds in length and serves as an introduction or a brief to establish the identity of the album. f) **Genuine**: a musical piece that deviates significantly from an artist’s typical style or the norm within a genre, exhibiting unique characteristics. These outliers are not classified as errors, noise, or other non-musical categories.

## 5.2 Automation Outlier Detection Result

The outlier detection algorithm was applied to the dataset, and the results obtained were analyzed to assess the performance of the method. The table 1 summarizes the performance of the outlier detection for various artists in the dataset. **True Positive Rate (TPR)** and **False Positive Rate (FPR)** represent the proportion of outliers that were correctly and incorrectly identified as outliers, respectively. **True Negative Rate (TNR)** and **False Negative Rate (FNR)** represent the proportion of non-outliers that were correctly and incorrectly identified as non-outliers, respectively. **Not Applicable (N/A)**: indicates no outliers being identified for a particular artist.

# 6 Discussion

## 6.1 Outlier Categorisation

Genuine outliers are complete songs that adhere to an artist’s typical musical structure but differ in sound and style. Often custom-made for specific events, they incorporate unique elements like distinct percussion. For instance, Colin Meloy’s “Lazy Little Ada” is situated at the center of Blue Oyster Cult’s cluster (Figure 4), deviating from Meloy’s

Table 1: The Result of Automatic Outlier Detection

Artist Name	TPR	FPR	TNR	FNR
Zakk Wylde	0	0.455	0.545	1
Blue Oyster Cult	1	0	1	0
Biohazard	0.5	0.381	0.619	0.5
Yoko Ono	0.25	0.049	0.951	0.75
Wilson Pickett	1	0.316	0.684	0
Wild Strawberries	1	0.087	0.913	0
Whitney Houston	0.5	0	1	0.5
Warren Zevon	1	0.356	0.644	0
Voivod	0.833	0	1	0.167
Uncle Tupelo	0.4	0.217	0.783	0.6
Triangle	0.75	0	1	0.25
Tim Deluxe	1	1	0	0
Thin Lizzy	0.429	0.37	0.63	0.571
Thievery Corporation	0.667	0.299	0.701	0.333
The Zombies	1	0.378	0.622	0
The Weather Girls	1	0.344	0.656	0
The Waybacks	0.5	0.214	0.786	0.5
The Tramps	N/A	0.136	0.864	N/A
The Subhumans	N/A	0.133	0.867	N/A
The Skatalites	N/A	0.317	0.683	N/A
THERION	0	0.275	0.725	1
The Mutton Birds	N/A	0.269	0.731	N/A
The Mission	N/A	0.324	0.676	N/A
The Mamas & The Papas	1	0.206	0.794	0
The Juan MacLean	N/A	0.324	0.676	N/A
Zee Avi	N/A	0.231	0.769	N/A
MNEMIC	1	0.043	0.957	0
Rod Lee	1	0.167	0.833	0
Colin Meloy	0.333	0.062	0.938	0.667

usual style yet resembling Blue Oyster Cult’s. In comparison, Meloy’s cluster exhibits lower loudness, potentially due to less percussion. When percussion is added to Meloy’s outliers, they show similarities to Blue Oyster Cult’s cluster. Conversely, most of Blue Oyster Cult’s songs form a well-defined cluster, but 4 outliers incorporate synthesizers and lack percussion, producing similarity to Meloy’s cluster.

In contrast, non-genuine outliers can be categorized into four types: Error, Speech, Sound Effect, and Intro. 1) Error tracks result from feature extractor misinterpretations or exceptions, such as doubling tempo or missing audio features. 2) Speech is often used in albums for storytelling, setting narrative themes, or as interludes to create a narrative flow between musical pieces, e.g., Kendrick Lamar’s “To Pimp a Butterfly.” 3) Sound effects enhance the musical narrative, especially in concept albums like Pink Floyd’s “The Trial.” 4) Intros are short tracks that set the tone, introduce themes, or provide transitions between songs. However, these tracks do not follow typical musical structures (e.g., ABAB, ABAC), so we consider them non-genuine outliers.

## 6.2 Outlier Detection Result

The overall results of the outlier detection algorithm show that it is capable of identifying outliers in an artist’s body of work. The algorithm performs exceptionally well on some artists by correctly classifying all outliers and non-outliers. However, varying performance among artists suggests the need for refining the algorithm with additional constraints. In summary, we can obtain these insights from the following results:






Fig. 4: Outlier Overlap: Shared characteristics between two artists' outliers.

The consideration of distinctiveness in the algorithm has worked well on artists that primarily have one style, for example, Blue Oyster Cult, The Mamas & The Papas, and Wild Strawberries had a perfect TPR and TNR, indicating that the definition works well for this artist even without considering the constraints of whether it forms a complete song. Furthermore, we find the distinctiveness constraint effective in identifying genuine music outliers in some artists. For example, the algorithm isolated outliers in artists such as Blue Oyster Cult (TPR: 1.0, FPR: 0, TNR: 1, FNR: 0) and Warren Zevon (TPR: 1.0, FPR: 0.356, TNR: 0.644, FNR: 0).

However, the varied performances across different artists suggest that discarding the consideration of constraints 1. Forms A Complete Song and 2. Non-adherence may lead to the algorithm's inability to accurately determine whether an outlier is genuine, as it lacks the capacity to assess the music structure of input recordings. The reasons causing these varying performances can be summarized as follows:

**Mixed content:** Recordings may contain various non-musical elements, including intros with predominantly speech content (e.g., "Rod Intro" by Rod Lee), live recordings featuring audience applause or speech interactions with the audience (e.g., "Dracula's Daughter" by Colin Meloy), or studio chats consisting solely of speech (e.g., "The Way I Feel Inside / Studio Chat" by The Zombies). These non-musical segments may introduce noise and impact the algorithm's performance.

**Noise and artifacts:** The presence of noise, artifacts, or other non-musical elements within a recording might lead to it being classified as an outlier, even if it does not constitute a genuine outlier in terms of musical content. Such factors can interfere with the algorithm's capability to accurately assess a song's structure.

**Transitional pieces:** Some artists release tracks with transition pieces containing sound effects or ambient sounds, which can be difficult for the algorithm to categorize as genuine outliers. Notable examples include "Catalepsy I" by Voivod and "The Audio Injection" by MNEMIC.

### 6.3 Limitations and Future Work

Overall, the findings of this study highlight the importance of considering outliers and their potential impact on the audio characteristics of different genres and artists in music production. It suggests several avenues for future research, including:




Fig. 5: Automatic Outlier Detection Using  $k$ -means: Demonstrating the algorithm works well on single-style artists (a) but fails for multi-style artists (b).

**Integration of more music features:** Expanding the range of music features integrated into the algorithm can potentially improve its performance in detecting genuine outliers. The current features, such as loudness and tempo, may not fully capture the characteristics and the style of music. Therefore, incorporating additional features like timbre, harmony, and chroma could enhance the algorithm’s effectiveness in distinguishing genuine outliers from the rest of an artist’s work.

**Consideration of “Forms A Complete Song” and “Non-Adherence” constraints:** Incorporating the “Forms A Complete Song” constraint ensures that the detected outliers are actual pieces of music. This guarantees that the outliers are indeed genuine musical outliers and not artifacts or other irrelevant audio content. The ‘Non-Adherence’ constraint ensures detected outliers distinctly deviate from an artist’s typical musical structure. This helps to identify unique songs that stand out from an artist’s typical style. Furthermore, music segmentation techniques can be considered to extract musical parts from audio pieces that may contain both speech and music. This will aid in ensuring the detected item actually is music.

**Handling artists with more than one style:** The current outlier detection approach may struggle with artists exhibiting multiple styles, such as Rod Lee (shown in fig.5). Addressing this limitation would result in a more representative understanding of an artist’s work and improve the accuracy of outlier detection. One possible solution is to consider an artist’s stylistic diversity when detecting outliers, thereby accounting for the various styles present within their body of work.

**Exploration of other clustering models:** In this study, only the  $k$ -means clustering algorithm was considered for the clustering model. However,  $k$ -means is based on circular data, which can lead to suboptimal results. Exploring other clustering models that may better handle the complexities of music data could further enhance the performance of the outlier detection algorithm and yield more accurate results.

## 7 Conclusion

In conclusion, this study has proposed a definition for genuine music outliers and explored the application of an outlier detection algorithm in music genre datasets. The

results have demonstrated that the consideration of distinctiveness is a reasonable starting point for detecting music outliers. However, the current approach lacks the ability to detect the music structure and struggles when handling artists with more than one style. To overcome these limitations, future work should focus on integrating more music features, such as timbre, harmony, and chroma, and considering the constraints of “Forms A Complete Song” and “Non-Adherence.” Furthermore, music segmentation techniques should be explored to extract musical parts from audio pieces containing both speech and music. Handling artists with multiple styles and exploring alternative clustering models, such as those that can better accommodate non-circular data, are other avenues for improvement. By addressing these limitations and incorporating these suggestions, the proposed outlier detection approach can be further refined and made more robust for detecting genuine music outliers in diverse music genre datasets.

## References

- [1] Ankerst, M., Breunig, M.M., Kriegel, H.P., Sander, J.: OPTICS. *ACM SIGMOD Record* **28**(2), 49–60 (Jun 1999). <https://doi.org/10.1145/304181.304187>
- [2] Ariyaratne, H.B., Zhang, D.: A novel automatic hierarchical approach to music genre classification. In: 2012 IEEE International Conference on Multimedia and Expo Workshops. IEEE (Jul 2012). <https://doi.org/10.1109/icmew.2012.104>
- [3] Arthur, D., Vassilvitskii, S.: How slow is the k-means method? In: Proceedings of the Twenty-Second Annual Symposium on Computational Geometry. p. 144–153. SCG '06, Association for Computing Machinery, New York, NY, USA (2006). <https://doi.org/10.1145/1137856.1137880>
- [4] Azcarraga, A., Flores, F.K.: A study on self-organizing maps and k-means clustering on a music genre dataset. In: Theory and Practice of Computation. WORLD SCIENTIFIC (Oct 2017). [https://doi.org/10.1142/9789813234079\\_0017](https://doi.org/10.1142/9789813234079_0017)
- [5] Bello, J.P.: Measuring structural similarity in music. *IEEE Transactions on Audio, Speech, and Language Processing* **19**(7), 2013–2025 (2011). <https://doi.org/10.1109/TASL.2011.2108287>
- [6] Bertin-Mahieux, T., Ellis, D.P.W., Whitman, B., Lamere, P.: The million song dataset. In: International Society for Music Information Retrieval Conference (2011)
- [7] Bountouridis, D., Koops, H.V., Wiering, F., Veltkamp, R.C.: Music outlier detection using multiple sequence alignment and independent ensembles. In: Amsaleg, L., Houle, M.E., Schubert, E. (eds.) *Similarity Search and Applications*. pp. 286–300. Springer International Publishing, Cham (2016)
- [8] Dai, S., Yu, H., Dannenberg, R.B.: What is missing in deep music generation? a study of repetition and structure in popular music. In: What is missing in deep music generation? A study of repetition and structure in popular music. arXiv (2022). <https://doi.org/10.48550/ARXIV.2209.00182>
- [9] Deng, D.: Dbscan clustering algorithm based on density. In: 2020 7th International Forum on Electrical Engineering and Automation (IFEEA). pp. 949–953 (2020). <https://doi.org/10.1109/IFEEA51475.2020.00199>
- [10] Fathollahi, M.S., Razzazi, F.: Music similarity measurement and recommendation system using convolutional neural networks. *International Journal of Multimedia Information Retrieval* **10**(1), 43–53 (Mar 2021). <https://doi.org/10.1007/s13735-021-00206-5>
- [11] Ghosal, S.S., Sarkar, I.: Novel approach to music genre classification using clustering augmented learning method (calm). In: AAAI Spring Symposium Combining Machine Learning with Knowledge Engineering (2020)

- [12] Hawkins, D.M.: Identification of Outliers. Springer Netherlands (1980). <https://doi.org/10.1007/978-94-015-3994-4>
- [13] Herskind Sejr, J., Christiansen, T., Dvinge, N., Hougesen, D., Schneider-Kamp, P., Zimek, A.: Outlier detection with explanations on music streaming data: A case study with danmark music group ltd. *Applied Sciences* **11**(5) (2021). <https://doi.org/10.3390/app11052270>
- [14] Knees, P., Schedl, M.: A survey of music similarity and recommendation from music context data. *ACM Transactions on Multimedia Computing, Communications, and Applications* **10**(1), 1–21 (Dec 2013). <https://doi.org/10.1145/2542205.2542206>
- [15] Lu, Y.C., Wu, C.W., Lerch, A., Lu, C.T.: Automatic outlier detection in music genre datasets. In: *International Society for Music Information Retrieval Conference* (2016)
- [16] Mehr, S.A., Singh, M., Knox, D., Ketter, D.M., Pickens-Jones, D., Atwood, S., Lucas, C., Jacoby, N., Egner, A.A., Hopkins, E.J., Howard, R.M., Hartshorne, J.K., Jennings, M.V., Simson, J., Bainbridge, C.M., Pinker, S., O'Donnell, T.J., Krasnow, M.M., Glowacki, L.: Universality and diversity in human song. *Science* **366**(6468) (Nov 2019). <https://doi.org/10.1126/science.aax0868>
- [17] Müller, M.: *Music Structure Analysis*, pp. 167–236. Springer International Publishing, Cham (2015). [https://doi.org/10.1007/978-3-319-21945-5\\_4](https://doi.org/10.1007/978-3-319-21945-5_4)
- [18] Neubarth, K., Conklin, D.: Identification and description of outliers in the densmore collection of native american music. *Applied Sciences* **9**(3) (2019). <https://doi.org/10.3390/app9030552>
- [19] Ng, R., Han, J.: Clarans: a method for clustering objects for spatial data mining. *IEEE Transactions on Knowledge and Data Engineering* **14**(5), 1003–1016 (2002). <https://doi.org/10.1109/TKDE.2002.1033770>
- [20] Panteli, M., Benetos, E., Dixon, S.: A computational study on outliers in world music. *PLOS ONE* **12**(12), e0189399 (Dec 2017). <https://doi.org/10.1371/journal.pone.0189399>
- [21] Paulus, J., Klapuri, A.: Music structure analysis by finding repeated parts. In: *Proceedings of the 1st ACM workshop on Audio and music computing multimedia*. ACM (Oct 2006). <https://doi.org/10.1145/1178723.1178733>
- [22] Sargent, G., Bimbot, F., Vincent, E.: A regularity-constrained Viterbi algorithm and its application to the structural segmentation of songs. In: *International Society for Music Information Retrieval Conference (ISMIR)*. Miami, United States (Oct 2011)
- [23] Schedl, M., Gómez, E., Urbano, J.: Music information retrieval: Recent developments and applications. *Foundations and Trends® in Information Retrieval* **8**(2-3), 127–261 (2014). <https://doi.org/10.1561/15000000042>
- [24] Shingte, G., d'Aquin, M.: Unsupervised learning approach for identifying sub-genres in music scores. In: *Irish Conference on Artificial Intelligence and Cognitive Science* (2019)
- [25] Singh, R., Nakamura, E.: Dynamic cluster structure and predictive modelling of music creation style distributions. *Royal Society Open Science* **9**(11) (Nov 2022). <https://doi.org/10.1098/rsos.220516>
- [26] Smiti, A.: A critical overview of outlier detection methods. *Computer Science Review* **38**, 100306 (Nov 2020). <https://doi.org/10.1016/j.cosrev.2020.100306>
- [27] Wang, J.C., Hung, Y.N., Smith, J.B.L.: To catch a chorus, verse, intro, or anything else: Analyzing a song with structural functions. In: *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE (May 2022). <https://doi.org/10.1109/icassp43922.2022.9747252>
- [28] Zhang, L., Jiang, F.: Visualizing symbolic music via textualization: An empirical study on chinese traditional folk music. In: *Mobile Multimedia Communications*, pp. 647–662. Springer International Publishing (2021). [https://doi.org/10.1007/978-3-030-89814-4\\_47](https://doi.org/10.1007/978-3-030-89814-4_47)