MIDI Music Ranking and Retrieval

Devlin Gallagher

Rochester Institute of Technology, Rochester NY 14623, USA erg4321@rit.edu
https://www.rit.edu/computing/study/computer-science-bs

Abstract. Algorithmic music recommendations are gaining popularity as streaming services employ various techniques to suggest music to users. However, MIDI files hold much more implicit information than an MP3, providing more objective comparators. The MIDI Music Ranking and Retrieval Program (MMRR) calculates ranking based on instrument similarity and note progression similarity. Based on a query MIDI file, the rest of the collection is then compared and ranked, and the results are returned in order, given the precision and recall parameters. This program combines aspects of retrieval and recommendation algorithms, revealing a new approach to music recommendation by using information retrieval techniques. The MMRR Program can build an overall ranking in a user's personal music collection, constantly refining a music queue by taking previous rankings into account without the use of streaming services. The program shows potential for generative contexts as well by using the rankings to determine the most likely note progressions given an excerpt of a piece.

Keywords: Information Retrieval · MIDI · Musical Similarity · Music Retrieval Systems

1 Information Retrieval Contexts and Applications

Information Retrieval (IR) is the process of finding relevant information from a collection of documents in response to a user's query. Traditional IR systems operate on text documents by indexing keywords, computing similarity measures, and ranking search results. The same principles can be extended to music, where each musical composition can be viewed as a document containing structured information instead of words.

Unlike compressed audio formats such as MP3 files, MIDI files store layered representations of music. These representations include note pitches, note durations, timing, tempo, instruments, and other information. Because these elements correspond directly to concepts in music theory, MIDI provides additional meaningful data for retrieval systems based on musical similarity rather than wave behavior.

1.1 Information Retrieval in Music

The MMRR system applies classic information retrieval concepts to a collection of MIDI files. Each MIDI composition serves as a document, while a query MIDI file functions as the search query. Instead of matching words or phrases, the retrieval process compares musical features extracted from each composition. The ranking algorithm considers two primary features:

- Instrument Similarity: a measurement of how closely two pieces match in instrumentation.
- Note Progression Similarity: a measurement of comparison in melodic sequences between compositions.

The resulting similarity score allows every piece in the collection to be ranked according to its relevance to the query. Standard IR evaluation metrics such as precision and recall can then measure how effectively the retrieval system identifies musically related pieces.

2 Role of Music Theory

Music theory provides the semantic foundation for the retrieval process. Musical compositions are defined by relationships among notes, rhythm, harmony, tempo, and instrumentation. Rather than treating music as raw audio data, the MMRR system compares symbolic musical structures.

The accompanying Python program demonstrates how synthetic MIDI files can be generated for experimentation. The program randomly selects:

- tempo ranges corresponding to traditional tempo markings such as Largo, Moderato, and Presto;
- MIDI instrument programs;
- note durations distributed across the length of the composition; and
- note pitches drawn from predefined musical note values.

Each generated MIDI file contains structured musical information that can later be indexed and compared. This synthetic dataset allows retrieval algorithms to be tested under controlled conditions before applying them to real musical collections.

3 Retrieval and Ranking Algorithm

The MMRR framework follows a typical information retrieval pipeline:

1. Generate or import a collection of MIDI files.
2. Extract musical features including instrumentation, tempo, and note sequences.
3. Represent each piece using these extracted features.
4. Compare a query MIDI file against every document in the collection.
5. Compute similarity scores based on instrument overlap and note progression.
6. Rank the collection according to similarity.
7. Return the highest-ranked musical pieces while evaluating performance using precision and recall.

3.1 Mathematical Formula

Given a query MIDI file Q and a collection of MIDI documents $\mathcal{D} = \{D_1, D_2, \ldots, D_n\}$, the MMRR algorithm computes a similarity score for every document in the collection. Each MIDI file is represented by two primary feature sets: its instrumentation and its note progression. The overall ranking score is defined as a weighted linear combination of these two similarity measures,

$$S(Q, D_i) = \alpha S_I(Q, D_i) + \beta S_N(Q, D_i), \tag{1}$$

where $\alpha, \beta \in [0, 1]$ and $\alpha + \beta = 1$. The parameters α and β determine the relative importance of instrumentation and melodic similarity.

Instrumentation is represented as a set of MIDI program numbers. Let I_Q and I_{D_i} denote the instrument sets of the query and document, respectively. Instrument similarity is computed using the Jaccard coefficient,

$$S_I(Q, D_i) = \frac{|I_Q \cap I_{D_i}|}{|I_Q \cup I_{D_i}|}. \tag{2}$$

The melodic component is represented as an ordered sequence of MIDI note numbers,

$$N_Q = (n_1, n_2, \ldots, n_k), \qquad N_{D_i} = (m_1, m_2, \ldots, m_\ell).$$

For two aligned sequences, note progression similarity is computed as

$$S_N(Q, D_i) = \frac{1}{\min(k, \ell)} \sum_{j=1}^{\min(k,\ell)} \delta(n_j, m_j), \tag{3}$$

where

$$\delta(a, b) = \begin{cases} 1, \&a = b, \\ 0, \&\text{otherwise}. \end{cases}$$

Once the similarity score has been computed for every document, the collection is ranked according to

$$R = \operatorname{argsort}_{D_i \in \mathcal{D}} \left(S(Q, D_i) \right), \tag{4}$$

where the resulting ordered list R contains the MIDI files from most to least similar. The effectiveness of the retrieval algorithm is then measured using the standard information retrieval metrics of precision and recall,

$$\text{Precision} = \frac{|\text{Relevant} \cap \text{Retrieved}|}{|\text{Retrieved}|}, \tag{5}$$

$$\text{Recall} = \frac{|\text{Relevant} \cap \text{Retrieved}|}{|\text{Relevant}|}. \tag{6}$$

These metrics quantify both the accuracy and completeness of the ranking generated by the MMRR algorithm.

This approach relies entirely on the musical content itself; and consequently, recommendations remain objective based on the user's independent collection.

4 Potential Applications

Content-based MIDI retrieval has several important applications beyond music recommendation.

For personal music libraries, the system can organize and recommend compositions without requiring an internet connection or streaming platform. Musicians and composers can search for pieces with similar instrumentation or melodic structure, making retrieval more meaningful than searching by title or artist alone.

The ranking methodology also supports generative music applications. By identifying the most similar note progressions across a collection, the system can estimate likely continuations of incomplete melodies. This creates opportunities for algorithmic composition, music completion, and intelligent accompaniment systems.

5 Conclusion

The MIDI Music Ranking and Retrieval program demonstrates how information retrieval techniques can be successfully combined with music theory to create a content-based recommendation system. By representing music through symbolic MIDI data rather than compressed audio, the system compares compositions according to their intrinsic musical characteristics, including instrumentation and melodic progression. This approach offers an interpretable alternative to traditional recommendation systems while providing additional applications in music organization, retrieval, and generative composition. As symbolic music analysis continues to evolve, integrating information retrieval with music theory has the potential to improve both music recommendation and computational creativity.