

Melodic Similarity and Applications Using Biologically-Inspired Techniques

*

Dimitrios Bountouridis*, Daniel G. Brown[†], Frans Wiering*, Remco C. Veltkamp[†]

*Department of Information and Computing Sciences, Utrecht University

[†]David R. Cheriton School of Computer Science, University of Waterloo

d.bountouridis@uu.nl, dan.brown@uwaterloo.ca, {f.wiering, r.c.veltkamp}@uu.nl

ABSTRACT

Melodic similarity is a complex concept that manifests itself in a number of Music Information Retrieval (MIR) tasks such as query-by-humming and cover song detection. Typically, similarity models are based on intuition or heuristics; thus, applicability to broader contexts cannot be guaranteed. We argue that data-driven tools and analysis methods, applied to melodies known to be related, can provide us with information regarding the fine-grained nature of music similarity. Interestingly, melodic and biological sequences share a number of parallel concepts; from the natural sequence-representation, to their mechanisms of generating variations, i.e., oral transmission and evolution respectively. As such, there is a great potential for applying data-driven scientific methods and tools from bioinformatics to music. Our paper relies on such methods to *a*) acquire new knowledge through a melodic stability analysis and *b*) model global melodic similarity and apply it to a retrieval/classification scenario.

ACM Reference Format:

Dimitrios Bountouridis[1], Daniel G. Brown[2], Frans Wiering[1], Remco C. Veltkamp[2]. 2018. Melodic Similarity and Applications Using Biologically-Inspired Techniques. In *Proceedings of Dutch-Belgian Information Retrieval Workshop (DIR2018)*. ACM, New York, NY, USA, 1 page.

1 SUMMARY

The current digital age allows listeners to stream massive collections of music. In addition, the proliferation of music streaming services has raised the listeners' interest in the accompanying chords, the lyrics, the original versions of a cover, and many more scenarios that service providers cannot deal with manually. This development brings Music Information Retrieval (MIR) to the centre of attention. The field includes research about accurate and efficient computational methods, applied to various music retrieval and classification tasks.

Such tasks require us to build representations of previously seen classes (e.g., sets of covers of the same song), which can be only

compared to a query (e.g., a cover song whose original is unknown) by means of a meaningful music similarity function. A robust MIR system should model the fuzziness and uncertainty of the differences between two musical items perceived as similar. However, this "knowledge", the exact mechanics of perceived similarity, is still unknown or incomplete. This is not surprising considering music's inherently complex nature.

To overcome, or avoid addressing the aforementioned issues, many MIR approaches to similarity rely on cognition studies, expert heuristics, music theory or formalized models in general. However, all such approaches have limited explanatory power and fail to generalize. In contrast, this paper focuses on acquiring knowledge on music and melodic similarity in particular from the data itself. Since data-driven methods and tools have been under development for years in bioinformatics, and since biological and music sequence share resembling concepts, we investigate their applicability inside a musical context.

First, this paper tackles the concept of meaningful and musically significant alignments of related melodies, by applying the bioinformatics structural alignment metaphor to music motifs. Our results reveal that the MAFFT multiple alignment algorithm, which uses gap-free sections as anchor points, is a natural fit for multiple melodic sequences; a strong indication of the importance of musical patterns for melodic similarity. Trusted alignments using MAFFT allow to organize melodic variations such that melodic stability/variation can be analysed. We therefore present a stability analysis free of heuristics or biases that might have been introduced following other approaches.

Secondly, this paper investigates the modelling of global melodic similarity. We capture the probability of one note to be changed to another in a variation and create musically appropriate note-substitution scoring matrices for melodic alignment. We then put these matrices successfully to the test by designing retrieval and classification tasks. Our data-driven modelling of music similarity outperforms the naive ± 1 matrix, indicating that indeed some novel knowledge is captured. Additionally, we show that variations inside a melody can be an alternative source for modelling the similarity of variations among tune families or cliques of covers.

In general, we show that bioinformatics tools and methods can find successful application in music, to answer in a reliable, data-driven way a number of important, on-going questions in MIR. We argue that data-driven approaches, such as ours, constitute an ideal balance between the two occasionally contradicting goals of MIR, problem solving and knowledge acquisition.

^{*}The full version of this paper was published at Applied Sciences 2017, 7(12), 2017 and can be found at: <https://www.mdpi.com/2076-3417/7/12/1242>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

DIR2018, Nov 2018, Leiden, the Netherlands

© 2018 Copyright held by the owner/author(s).