Towards Logic-based Representations of Musical Harmony for Classification, Retrieval and Knowledge Discovery

Amélie Anglade AMELIE.ANGLADE@ELEC.QMUL.AC.UK
Simon Dixon SIMON.DIXON@ELEC.QMUL.AC.UK
Centre for Digital Music, Queen Mary University of London, Mile End Road, London E1 4NS, UK

Abstract

We present a logic-based framework using a relational description of musical data and logical inference for automatic characterisation of music. It is intended to be an alternative to the bag-of-frames approach for classification tasks but is also suitable for retrieval and musical knowledge discovery. We present the first results obtained with such a system using Inductive Logic Programming as inference method to characterise the Beatles and Real Book harmony. We conclude with a discussion of the knowledge representation problems we faced during these first tests.

1. Introduction

Most of the automatic music classification systems (performing e.g. genre classification) for audio data are based on the so-called "bag-of-frames" approach. In this approach, feature vectors (typically spectral features such as MFCC) are computed over short frames. Then a classifier computes the global distribution or average values of these vectors over the whole piece or passage for each class. However this statistical method presents several limitations. For instance it has been shown that contrary to the bag-of-frames assumption the contribution of a musical event to the perceptual similarity is not proportional to its statistical importance (rare musical events can even be the most informative ones) (Aucouturier et al., 2007). Moreover the bag-of-frames approach ignores temporal organisation of the acoustic signal. However when comparing pieces from similar genres or passages of a same song it is crucial in the retrieval process to use sequences and not only average values or global statistical distributions of features over a whole passage or piece (Casey & Slaney, 2006).

We believe a logic-based representation of the musical events together with logical inference such as Inductive Logic Programming (ILP) of higher level phenomena would overcome these limitations. Indeed temporal relations between musical events are easily expressible in a relational framework. Moreover logical inference of rules allows to take into account all events, even those which are rare. Another advantage of using logical rules is that they are human-readable.

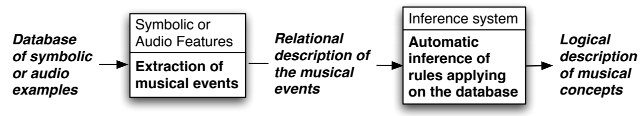

Figure 1. Logic-based framework for musical rule induction

Additionally logical inference has already been successfully used in several music-related projects, for instance to induce rules about popular music harmonisation (Ramirez, 2003), counterpoint (Morales & Morales, 1995) and expressive performance (Widmer, 2003).

2. The Reasoning Framework

Taking the example of these promising applications of logic to music we intend to build a logic-based reasoning system able to characterize songs for classification or similarity evaluation purposes. An illustration of this system is given in Figure 1. It takes a database of examples to characterise. Notice that these examples can either be audio signals or symbolic examples. These examples are then analysed by a musical event extractor using either symbolic or audio features. Finally the relational description of the examples resulting from this analysis is given to an inference system which derives musical rules that are true for the examples. An example of such an automatically derived rule for the blues genre would be: $\mathit{12\_bar\_structure(X)} \wedge \mathit{blues\_scale(X)} \wedge \mathit{syncopation(X)} \Rightarrow \mathit{blues(X)}$

We imagine three kinds of applications for this system:

1. **Classification and retrieval** of pieces similar to a given set of examples, for instance for recommender systems.

2. **Musical knowledge discovery**: some of the derived rules might describe new interesting musical phenomena.

3. **Query in musical databases**: finding music that matches a pattern/description expressed by the user in the form of a logic formula (either defined by hand or a rule previously derived by the system) so without any audio or score example.

3. First Results

We ran a first set of experiments using the ILP system Aleph (Srinivasan, 2003) which is based on Inverse Entailment. Our objective was to characterize the harmony of the 180 songs featured on the Beatles' studio albums (containing a total of 14,132 chords) and of 244 jazz standards from the Real Book (24,409 chords). More precisely we focused on characterising chord sequences of length 4 in these corpora. To avoid the event extraction step we analysed manually annotated chord data[1] containing information about the root, bass note and component intervals (relative to the root) of the chords. We pre-processed these data to obtain high-level information such as chord category (e.g. major, minor, suspended), degree (e.g. tonic, dominant) and intervals between chords, before passing them to an ILP system which extracted the harmony rules underlying them. It generated 3667 harmony rules characterising and differentiating the Real Book songs and the Beatles music. Examples of learned rules are given in Table 1. Encouragingly some of these rules cover well-known jazz or pop harmonic patterns such as ii-V-I-IV and the "turnaround" pattern I-VI-II-V (for the Real Book), and the I-IV-I-V pattern (for the Beatles). It was also possible to identify patterns specific to the Beatles such as the extensive use of major chords (the predominant maj-maj-maj-maj pattern) and the cyclic patterns (e.g. I-IV-V-I-IV...) characterising their early compositions. More details on these experiments can be found in (Anglade & Dixon, 2008).

4. Discussion

Even though these first results are encouraging we identified some knowledge representation problems arising when using ILP as inference technique. In ILP both the concept and the vocabulary to describe it (or background knowledge) needs to be defined in advance. Thus, two different vocabularies result in different descriptions of the concept. The process of concept characterisation is interactive: to end up with an interesting theory we might need to manually refine the vocabulary several times. For instance for the Beatles in our final attempt we used the concept "chord sequence (of length 4) in the Beatles" (`chord_seq/4`) and asked the system to describe it in terms of "chord category sequence" (`category_seq/8`), "root interval sequence" (`rootInterval_seq/7`), etc. But our earlier attempts at using only low level concepts (e.g. using independent descriptions of each chord, linked only by a "predecessor" predicate) failed to provide meaningful descriptions of the target concept. However it seems that as we refine the vocabulary we inevitably reduce the problem to a pattern matching task and not to a pattern discovery task as we intend to, allowing us to validate or refute hypotheses about the concept but not to make any really novel (knowledge) discoveries.

[1]RDF descriptions of the Real Book chords (available on http://chordtranscriptions.net/) and of Harte's transcription of the Beatles (available on request)

[Rule 16][Pos cover=7.86%]: chord_seq(A,B,C,D):-
rootInterval_seq(A,B,C,D,[perf,4th],[perf,4th],[perf,4th]).
[Rule 25][Pos cover=4.09%]: chord_seq(A,B,C,D):-
category_seq(A,B,C,D,min,dom,min,dom).
[Rule 57][Pos cover=2.01%]: chord_seq(A,B,C,D):-
bassInterval_seq(A,B,C,D,[maj,7th],[perf,4th],[perf,4th]).

Table 1. Examples of rules learned on the Real Book data

5. Acknowledgments

This work is part of the OMRAS2 project funded by EPSRC grant EP/E017614/1.

References

Anglade, A., & Dixon, S. (2008). Characterisation of harmony with Inductive Logic Programming. *Under review for ISMIR 2008.*

Aucouturier, J.-J., Defreville, B., & Pachet, F. (2007). The bag-of-frames approach to audio pattern recognition: A sufficient model for urban soundscapes but not for polyphonic music. *JASA*, *122*, 881–91.

Casey, M., & Slaney, M. (2006). The importance of sequences in musical similarity. *Proc. of ICASSP 2006.* Toulouse, France.

Morales, E., & Morales, R. (1995). Learning musical rules. *Proc. of the IJCAI-95 Int. Workshop on Artificial Intelligence and Music.* Montréal, Canada.

Ramirez, R. (2003). Inducing musical rules with ILP. *Proceedings of ICLP2003.* Springer-Verlag (LNCS).

Srinivasan, A. (2003). The Aleph Manual (version 4).

Widmer, G. (2003). Discovering simple rules in complex data: A meta-learning algorithm and some surprising musical discoveries. *AI*, *146*, 129–148.