MEANINGFUL MUSIC RETRIEVAL

Frans Wiering
Department of Information and Computing Sciences, Utrecht University, Netherlands
Frans.Wiering@cs.uu.nl

1. INTRODUCTION

Given the present paradigm of Music Information Retrieval (MIR) research, one could claim that we will have attained all research goals as soon as we can answer any musical query with 100% precision and 100% recall. But is this really what drives MIR research and what its use consists in? I do not think so. MIR researchers are generally driven by a notable curiosity about music (see for example [6]). And from a user's perspective, MIR is not so much about extracting a complete set of items of musical information as about experiencing music that is interesting, enjoyable and, to a certain extent, novel or surprising. The motivation for MIR is thus primarily a musical one. Successful MIR research and applications should therefore contribute to our understanding and experience of music. The principal reason for people to compose, listen to, study, play and buy music is that it is meaningful to them. So my message about the future of MIR is to explore and exploit musical meaning, and to create applications that are able to support the human process of musical meaning generation. Even more so, in my vision, human musical behaviour should become one of the foundations of the discipline.

2. THE PROBLEM

There is much in MIR nowadays that inhibits a deep involvement with musical meaning. Drawing a slightly exaggerated, somewhat one-sided picture, mainstream MIR research proceeds as follows. Pieces of music, whether audio or symbolically represented, are collected, labelled and stored as objects rather than as representations of temporal processes. These objects are dissected by means of segmentation and feature extraction. The musical knowledge involved is mainly of an 'anatomical' nature and is only occasionally critically evaluated before its application. The modelling and comparison of features is driven by technology, treating the features as 'information units' that can be manipulated as if they were musical keywords [9]. The ubiquitous criterion is 'similarity' despite the well-known fact that musical similarity is an 'ill defined' concept, at least from a computational point of view. Finally, the retrieval results are evaluated against a given human ground truth.

The notion of ground truth in particular illustrates my point. The function of the ground truth is to act as a domain-derived Golden Standard against which the performance of MIR methods is measured. The criterion is to what extent this given ground truth can be reproduced. The term 'given' is highly relevant here, for the musical knowledge has become fixed and is no longer open to critical inspection. In effect, a hermetic boundary exists between MIR and musical knowledge [4]. It enables a peaceful coexistence of the two disciplines, but prevents true interdisciplinarity. The answer is to focus the evaluation not just on performance but on gaining insight into the ground truth as well. This will obviously lead to a better understanding of music, which in turn will lead to better approaches for MIR. Willard McCarty's methodology for humanities computing, his modelling theory in particular, may serve as a guideline for systematic discovery of interdisciplinary knowledge [5].

Furthermore, we tend to treat ground truths as if they represent straightforward observations of musical content, and consider variation in judgements as the outcome of random processes. But the ground truths that are used for example in evaluating musical similarity are the result of very complex processes, in which perception, personal background, training and taste are at least as important as the musical content that is being judged. The important issue is then not to model the average judgment but to understand how the individual differences arise. Here, I argue, the study of musical meaning can make an important difference.

Several arguments could be adduced against taking a 'meaning' viewpoint in MIR, for example:

- *Musical meaning is too subjective.* I will discuss this issue in the next section.
- *It is all in the data.* This position underestimates the role of music cognition and tends to treat music as an object rather than as a process.
- *It is interesting, but no science.* Good music research is methodologically rigorous as well, and there are good models for bridging methodological gaps [5].
- *It is too specialist.* Musical meaning is not just relevant to specialists (though worldwide they are a huge market), but to everyone who cares about music.
- *It is not relevant to MIR applications.* Understanding the subjective process of meaning generation will help to create more focused models of musical similarity, allow advanced personalisation and generally

Permission to make digital or hard copies of all or part 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.

© 2009 International Society for Music Information Retrieval

contribute to the design of truly useful MIR applications.

3. MUSICAL MEANING

But what is musical meaning? One might seek an analogy in language. Language enables us to refer to and reason about things in the real world. Meaning resides in the relationship between these things and the concepts in our mind. Such a relationship is at best weak in music: music is not emphatically 'about' things and processes in the real world. But when we listen to music, our mind actively picks up regularities and manipulates musical structures that we non-consciously extract from sound [2]. This manipulation generates a neurophysiological response [7]. Such responses produce a meaningful sensation, especially if they are secondarily associated with aspects of the listener's physical and social environment. So musical meaning is not about concrete referential 'semantics' shared by a large population, but about a personal network of connotations that associatively connect music to emotions and the listener's experience of his/her environment.

The specific meaning that a piece of music has to someone is a personal matter by definition. Small wonder that musical meaning was regarded for a long time as an illegitimate question in music research. Yet questions about meaning do not just go away when they are being ignored, as they relate to the fundamental reasons why we want to study at all. Meaning has come back as a central topic in 'the new musicology,' where it is answered using a variety of methods. Some of these are highly relevant in a MIR context.

According to Lawrence Kramer, the new musicology is probably best described as cultural musicology [3]. Its aim is to gain an understanding of musical subjectivity in history. Musical meaning has a central place in this study. Subjectivity is not seen as the individual being a 'private monad,' but as an individual's 'disposition to engage in social and cultural practices.' Musical meaning emerges as the result of an action, a negotiation of musical text (in any suitable medium) and context. Meaning thus does not reside in the musical structure itself, but certain meanings are made possible by aspects of the musical structure. In a similar way, Nicholas Cook describes the engagement between the listeners and the 'potential meanings' in a musical piece, by which meaning is actualised [1].

Even though the cognitive and musicological approaches study musical meaning from a different point of view, there are interesting similarities:

- music has structural aspects that carry a potential for meaning generation;
- meaning is the outcome of a process that involves musical input and the disposition of the individual;
- the environment/context is an important factor in shaping the disposition of the individual.

Using these similarities one can get a grip on the subjectivity of musical meaning and open it up for MIR research and applications.

4. A MODEL FOR MEANING GENERATION

Figure 1 shows my high-level model of the generation of musical meaning in humans. A musical input stream is perceived and analysed into perceptual patterns. Perception itself is determined partly by biological constraints and partly by previous exposure, hence the arrow from the patterns and rules—the border between the two is vague—in long-term memory. In the next step (which actually may consist of a number of interrelating processes) the perceived pattern is interpreted on the basis of patterns and rules from long-term memory, and the representation of the current piece that resides in long-term memory (if it is a known piece) and is at the same time being built up in working memory. The interpreted pattern (or the process of interpretation) leads to a neurophysiological response that can be associated with the

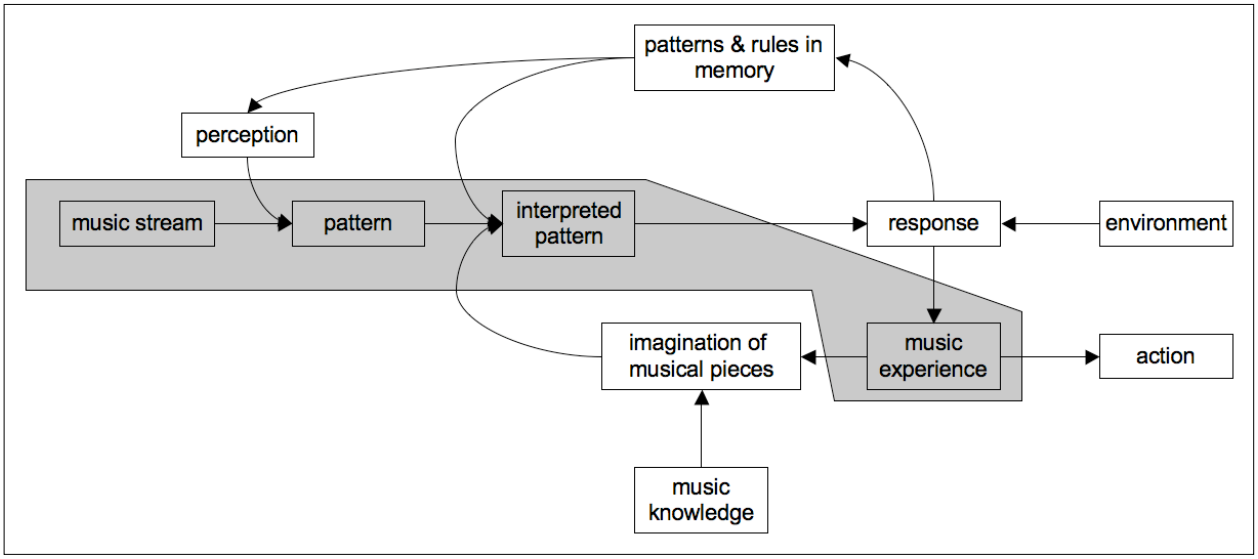

Figure 1. A model for the generation of musical meaning.

listener's present environment. During the process, the patterns and rules in the long-term memory may be updated. (This is certainly the case in infants and when being exposed to a new kind of music). The response also leads to a conscious experience of the music that may take various forms. This experience influences one's representation of the current piece and guides the interpretation of its continuation. It may also lead to a wide range of actions, such as moving or dancing to music, joining the performance, create new music, wanting to hear more such music, buying the music or, if the music belongs to a commercial, the article that is being advertised. Finally, one's representation of musical pieces is also shaped by knowledge of all sorts, for example biographical information, music theory and reviews.

The classical MIR approach is to predict the listeners' similarity experience on the basis of the input stream. This roughly corresponds to the shaded area in Figure 1. Note that the feedback loops in the model all fall outside of this area. Modelling as much of the process as feasible in MIR applications is an obvious goal to aim at. After all, modelling the musical mind of users of MIR systems will help to understand similarity, user needs, and enable sophisticated modes of personalisation. Also, interesting insights in music may be acquired. The following (conceptual) scenario may serve as a brief example.

Felix Mendelssohn Bartholdy (1809-1847) is probably one of the best-documented composers ever [8]. An incredible amount of personal papers have survived, including counterpoint exercises, sketches of compositions, letters, drawings and diaries of family members. Many figures in his direct surroundings are also well documented. As a consequence, we are extremely well informed about what music he was exposed to and when that happened. It would be an interesting challenge to try to reconstruct his musical mind, that is, his set of rules and patterns, and his representation of the musical works he knew, taking its development over time into account. This model would then provide the context from which his own works emerged. A dense network of relations between his works and his background could then be established, enhancing our understanding of his music. A simple MIR application of this model could direct people that are interested in a particular work by Mendelssohn to compositions by him and other composers that are related to this particular work in a meaningful way.

5. CONCLUSION

The scenario sketched in the previous section is of course completely unrealistic given the present state of research, and may require for its realisation knowledge that cannot conceivably be attained. Such scenarios, of which many more should be studied, can however help to formulate a research agenda of more feasible goals and to define concrete projects. Some possible items on this agenda are:

- creating computational models of working memory that treat a piece of music as a process in time;
- model the creation of rules and patterns in long-term memory;
- research musical similarity in order to develop more focused, personalisable models;
- investigate other types of musical relevance than similarity;
- study the process of meaning generation in ecologically valid experiments;
- design methods for assessing the contribution of MIR models to musical understanding;
- study new application possibilities enabled by meaningful music retrieval.

To draw up a comprehensive research agenda for meaningful music retrieval (MMR) seems an important goal for the fMIR workshop. The realisation of this agenda will keep researchers occupied for many years to come.

6. REFERENCES

[1] N. Cook. Theorizing musical meaning. *Music Theory Spectrum* 23, 170-95. 2001.

[2] I. Cross. Music, science and culture. *Proceedings of the British Academy* 147, 147-165. 2007

[3] L. Kramer. Musicology and meaning. *The Musical Times* 144 (1883), 6-12. 2003

[4] P. van Kranenburg, J. Garbers, A. Volk, F. Wiering, L.P Grijp & R.C. Veltkamp. Towards integration of MIR and folk song research. *Proceedings ISMIR 2007*, 505-508. Vienna: Austrian Computer Society, 2007.

[5] W. McCarty. *Humanities computing*. Palgrave Macmillan. 2005

[6] D. Meredith. "So why do we bother?": The goals of computational musicology and modelling music cognition. In: E. Selfridge-Field, F. Wiering, G. A. Wiggins (Eds.) *Knowledge representation for intelligent music processing. Dagstuhl Seminar 09051*. 2009. http://drops.dagstuhl.de/portals/index.php?semnr=09051

[7] I. Peretz. The nature of music from a biological perspective. *Cognition* 100, 1–32. 2006

[8] R. L. Todd. *Mendelssohn. A life in music.* Oxford University Press, 2003.

[9] F. Wiering. Can humans benefit from music information retrieval? In: S. Marchand-Maillet, E. Bruno, A. Nürnberger & M. Detyniecki (Eds.), *Adaptive Multimedia Retrieval: User, Context, and Feedback. 4th International Workshop, AMR 2006.* Lecture Notes in Computer Science 4398, 82-94. Springer 2007.