THE USEFULNESS OF QUALITATIVE THEORIES OF MUSICAL PERCEPTION

Gerhard WIDMER (gerhard@ai-vie.uucp)
Austrian Research Institute for Artificial Intelligence and
Department of Medical Cybernetics and Artificial Intelligence,
University of Vienna
Austria

ABSTRACT: The paper discusses the idea of developing and formulating *qualitative theories of (tonal) music* as a basis for intelligent musical computer systems. In particular, it concentrates on the use of such qualitative theories for programs that are to *learn* about music. It is argued that qualitative theories are a very natural formalism for describing intuitive, abstract knowledge about music, that such theories can greatly enhance the learning process in a learning system, and finally, that musicology may gain some interesting insights from experimental work involving AI systems and qualitative musical models.

1 Introduction

In a recent theoretical paper (Laske, 1988), O.Laske has emphasized the need for a new understanding of the discipline of 'musicology', broadening its scope to include both the construction of cognitive models of musical competence and performance, and empirically testing the adequacy of such models in working computer programs. We fully agree with this view, and propose that *qualitative models of musical perception* should be one of the central concerns of what Laske calls *theoretical musicology*. Moreover, we have developed a general qualitative theory (model) of the perception of a simple subdomain of tonal music, namely, harmonized melodies. We have also implemented a working system that can *learn* to harmonize given melodies, based on this qualitative perception model. Experimental results show that the availability of such a model greatly increases the effectiveness and efficiency of the learning process. The details of the model and the learning system are described in (Widmer, 1990b) and (Widmer, 1990c), respectively.

The current paper cannot describe these projects in detail; instead, it will provide some general arguments for the necessity and adequacy of qualitative theories of musical perception 1) by arguing that there *is* indeed common, 'intuitive', abstract musical knowledge shared by most listeners that can be modelled in a general way, 2) by suggesting that techniques and formalisms from *qualitative modelling* (a sub-field of Artificial Intelligence) are highly adquate for this task, and 3) by hinting at the various ways in which such a qualitative model can be used by an intelligent musical system. We will also argue that musicology can gain very interesting insights from experimental work involving the formulation of qualitative theories and their testing in real, practical applications.

2 Basic musical knowledge

The long-term goal of our research is the development of computer programs that 'understand' (tonal) music in a similar way as experienced listeners understand or perceive it. These systems will able be to *reason* about music, to *communicate* their deliberations to musicians in intelligible, 'common–sense' terms, and to *flexibly adapt* to new tasks and *learn* new musical concepts effectively. Obviously, all these tasks require a lot of music–specific knowledge. In order to keep the approach as general as possible, our project heavily relies on methods and techniques from the domain of *Machine Learning*. The idea is to give a learning program all the knowledge about the world of music that is readily available and can easily be formulated, and let the program learn the details by itself. For this learning scenario not to look like cheating, we must require the *a priori* knowledge given to the system to be some kind of 'natural', 'intuitive', in some ways 'basic' knowledge – the kind of knowledge that people can be assumed to possess also.

Here an interesting question arises: is there general musical knowledge that could be regarded as a psychologically and musicologically plausible basis? We believe that there is indeed such basic musical 'knowledge' (even if it is not explicit or conscious) – it is the sum total of what might be called 'habits of perception', the way people familiar with tonal music perceive and organize music in their minds (see, e.g., Lerdahl & Jackendoff, 1983). It is the kind of 'knowledge' that ordinary listeners in our culture have acquired from years of exposure to tonal music and that enables them to make 'intuitive' judgments such as whether a musical passage sounds good or bad, interesting or boring, coherent or incoherent, 'logical' or 'surprising', and the like. The term 'knowledge' is not meant to imply that listeners are necessarily aware of the factors that influence their decisions, let alone that they dispose of a vocabulary to explain their intuitions. What matters is that they have somehow learned to organize what they hear in mental structures that give coherence and 'meaning' to the music. We sum up all these types of 'knowledge' under the heading 'Perception of tonal music'. In the first stage of our project, we have identified a number of specific and abstract factors influencing the perception of harmonized melodies, and on this basis, have implemented an approximate, *qualitative* model that allows the computer to draw plausible inferences about a given piece.

3 Qualitative theories and models

A *qualitative model or theory* of some problem domain is some model that specifies the most important parameters/entities/concepts of the domain and the most important relationships among them, but only in a *qualitative* (as opposed to *quantitative*, precise) way. A qualitative model is by necessity an *abstraction* of the real world; in particular, it may

- be *too abstract* in at least two ways:
  - it does not deal with precise quantitative values; quantities are usually described only through qualitative terms like *high*, *low*, *increasing*, *decreasing*, etc.
  - it specifies mainly dependencies between different parameters, but not exact relationships; a typical example are Forbus' *qualitative proportionalities* (Forbus, 1984).
- be incomplete (it misses some relevant concepts and/or relationships)
- be incorrect (the model contains statements that are not true in the 'real world')
- refer to observable as well as to non–observable variables/concepts

Qualitative models (and formalisms for such models) developed so far were mainly intended to describe physical systems and processes; the most prominent approaches are Qualitative Process Theory (Forbus, 1984), QSIM (Kuipers, 1986) and DeKleer's theory of confluences (DeKleer & Brown, 1984). However, their emphasis on approximate, qualitative values and their ability to express abstract relationships between different parameters also make them the formalism of choice for models of music perception.

The perception model for harmonized melodies we have developed centers on the notion of an *effect*. Effects are sensations evoked in the listener by some features of a musical event. Such effects can arise on different *levels of abstraction*. For instance, two successive musical events that are very different in some respect may create a feeling of *contrast*; harmonic progressions may be felt to *smooth* or *abrupt*, and the like. On a more abstract level, *contrast*, *diversity*, and *buildup of tension* are factors that can contribute to an overall feeling of *interestingness* of a passage. Our model specifies perceivable effects on various hierachical levels of abstraction, and also describes, in a qualitative way, *influences* between effects (for instance, whether the presence of one effect strengthens or weakens the perceived prominence of some other effect).

For instance, the following qualitative statement describes the relation between the smoothness of the transition between adjacent chords and the perceived coherence of the musical passage:

QUAL-POS-PROP (smoothness-of-transition(C1, C2), coherence(C1, C2))
("perceived smoothness of transition from a chord C1 to a chord C2 is approximately positively proportional to perceived coherence of this musical passage; i.e., the smoother a chord progression sounds (which again depends on the relative harmonic distance between the chords), the more coherent will this passage be perceived")

In addition, the model takes into account relative *metrical* and/or *structural salience* of musical events, based on a plausible metrical analysis and time-span reduction (see Lerdahl & Jackendoff, 1983). These factors, too, can influence the perceived prominence of audible effects. Taken in its entirety, the model describes plausible hearings of harmonized melodies by suggesting various effects that a particular musical passage might create. For a detailed description of the model, see (Widmer, 1990b).

4 The usefulness of qualitative theories of perception

The model can be used by a computer program in a variety of ways to make intelligent inferences and judgments. However, because of the inherent incompleteness and abstactness of the model, all of these inferences will, at best, be *plausible inferences* (see Collins & Michalski, 1989).

- *Explanation, verification, prediction:* Given some musical passage, the system can explain or verify that the passage has certain qualities, by traversing its model top-down and verifying that certain conditions hold. Likewise, by traversing the model from bottom to top, the system can interpret a given piece of music in the sense of assessing its qualities and predicting how listeners will probably perceive it.
- *Plausible generalization, learning:* In situations where the system is trying to learn new musical concepts or induce general rules from observations (see, e.g., Michalski 1983), the model, though maybe too abstract and not entirely pertinent to the concepts to be learned, can provide strong hints concerning the *relative plausibility* of possible generalizations, so that the system can concentrate on those hypotheses that seem to make more musical sense, and thus learn faster and more effectively (Widmer 1990a,c).
- *Problem solving:* Though qualitative models are not well suited for direct use in problem solving, they can be extremely useful in the (semi-)automatic construction of specialized problem solvers for musical tasks by supporting effective learning processes: given such basic musical knowledge, a learning system can communicate with human experts, watch their problem solving behaviour, and learn specific rules for problem solving by plausible generalization (see above).

Apart from experiments of theoretical interest, one can imagine a variety of possible applications in practical musical systems. For instance, the explanatory capability of the model (or of a more complex model that covers a larger sub–domain of tonal music) could be exploited in intelligent musical assistants to composers, in systems employed in music education, etc. In all these scenarios, a qualitative model of perception would enable the system to evaluate given musical pieces, create its own compositions and rate their plausible quality, and explain its decisions and solutions to human experts in intelligible terms.

5 Possible benefits of such projects for musicology

Turning to more general issues, it seems clear that music theory and the cognitive psychology of music can reap some important benefits from AI experiments with qualitative theories of musical perception and organization. To take just one, we can see how experiments combining such qualitative theories and learning programs might shed some new light on the adequacy of the underlying musical assumptions: if the rules of a particular musical style are easily learnable on the basis of some abstract model, there must be a natural connection between the assumptions about musical perception made in the model and the specific rule system that defines the musical style. In such a way, different qualitative theories of musical perception could be tested in relation to the stylistic conventions of various musical epochs. This is an instance of the kind of interaction Laske envisioned between theoretical and applied musicology: practical musical system can be used to verify the adequacy (or inadequacy) of general musical competence theories. There might be a lot of explanatory potential in such experiments.

References

Collins, A. and Michalski, R. (1989). The Logic of Plausible Reasoning: A Core Theory. *Cognitive Science 13*, pp.1–49.

de Kleer J. and Brown J.S. (1984). A Qualitative Physics Based on Confluences. *Artificial Intelligence 24(1–3)*, pp.7–84.

Forbus, K.D. (1984). Qualitative Process Theory. *Artificial Intelligence 24(1–3)*, 85–169.

Kuipers, B.J. (1986). Qualitative Simulation: *Artificial Intelligence 29(3)*, 289–338.

Laske, O. (1988). Introduction to Cognitive Musicology. *Computer Music Journal 12(1)*, 43–57.

Lerdahl, F., and Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. Cambridge, Mass.: MIT Press.

Michalski, R. (1983). A Theory and Methodology of Inductive Learning. In R.Michalski, J.Carbonell, T.Mitchell, eds., *Machine Learning: An Artificial Intelligence Approach*. Palo Alto, CA: Tioga.

Widmer, G. (1990a). Learning a Complex Musical Task on the Basis of a Plausible Theory of Musical Perception. In *Proceedings of the ECAI Workshop on Artificial Intelligence and Music*, ECAI–90, Stockholm.

Widmer, G. (1990b). *A Qualitative Model of the Perception of Melodies and Harmonizations*. Submitted. Available as Report TR–90–12, Austrian Research Institute for Artificial Intelligence, Vienna.

Widmer, G. (1990c). *Using a Qualitative Perception Model to Learn Harmonization Rules for Melodies*. Submitted. Available as Report TR–90–13, Austrian Research Institute for Artificial Intelligence, Vienna.