

AHRC ICT Methods Network Workshop

THE FUTURE OF INFORMATION TECHNOLOGY IN MUSIC RESEARCH AND PRACTICE

GOLDSMITH'S COLLEGE, UNIVERSITY OF LONDON, 8 SEPTEMBER 2006

Report by David Meredith

Introduction

This workshop was proposed as a follow-up to the first AHRC ICT Methods Network Expert Seminar on Music, entitled 'Modern Methods for Musicology: Prospects, Proposals and Realities', chaired by Tim Crawford and held at Royal Holloway on Friday 3 March 2006.

Several important issues and themes emerged from the first expert seminar on music. The first of these was the need for a robust technological infrastructure for music-related ICT, including architectures, protocols and representations that support the development of flexible, extensible, affordable and interoperable music processing systems. Much discussion also focused on the problems involved in designing music software systems that are both powerful and easy to use. Another question that stimulated considerable debate at the expert seminar was whether ICT will or should cause a gradual evolution or a sudden revolution in musical practices. One of the main conclusions reached was that there is an urgent need to raise trans-disciplinary awareness in the field: music specialists should be made more aware of the limitations and potentials of technology; and technologists should better understand the real needs of music practitioners. It was also generally agreed that considerable effort should be put into promoting a culture of inter-disciplinary collaboration in which we exploit rather than fear the knowledge of those who are expert in fields other than our own.

It was hoped that the workshop reported on here would provide an opportunity for these issues to be discussed in more depth and thereby make a significant contribution to raising trans-disciplinary awareness and promoting a culture of inter-disciplinary collaboration. The principal aim of this workshop was therefore to allow experts in music with an interest in technology to talk to experts in technology with an interest in music and identify ways in which they can collaborate fruitfully to achieve worthwhile goals.

The workshop consisted of four 45-minute sessions, each with two discussions running in parallel. These sessions were interspersed with three, half-hour coffee breaks and an hour-long buffet lunch during which participants were able to continue their discussions and follow-up on issues in smaller groups or on a one-to-one basis. The discussions during these breaks were very animated and several participants commented afterwards on how valuable these long breaks had been for further exploration of possible avenues for collaboration with people they had met for the first time at the event. At the end of the day, each discussion chair presented a 15-minute summary of his or her discussion to all the participants.

The following sections summarize the discussion sessions.

Connecting the two cultures: How empirical does musicology need to be?

This session was proposed and chaired by Tim Crawford (Goldsmiths College, University of London). In his introduction, Crawford observed that the German word 'Musikwissenschaft' exemplifies an attitude to (or at least an aspiration of) musicology (not very fashionable in the UK these days), that it is (or can be) in some sense a scientific discipline. He noted that the multi-disciplinary process of developing intelligent and useful ICT methods and tools throws up a number of issues arising from the mismatch between the 'scientific

method' and humanistic scholarship which could affect the outcomes of such work in positive or negative ways. Crawford suggested that the discipline of engineering in some sense steers a middle course between these extremes and may offer a means for collaboration.

The discussion that ensued focused on the relationships between musicology, science and e-Science. Music is both an art and a science. Engineering perhaps takes a middle course between these 'extremes', because it is motivated by the need for practical solutions; while this may not be satisfactory from a pure-science point of view it is attractive for the musicologist as a 'client'.

Maybe we have a more urgent need for fast ways to get imperfect results than systematic, exhaustive methods. This may be a caricature of the scientific method, but the pursuit of pure knowledge is not necessarily the most efficient route to developing useful tools. Consider, for example, the Google search engine, which is based on a certain amount of science, plus some good engineering, but which is above all available and useful, albeit flawed for most scholarly purposes.

Typically, an academic musicologist has to persuade his or her peers using arguments that are based on imperfect or ambiguous evidence. However, science typically offers more 'precise' solutions and, as the problem of 'expert witnesses' in court reminds us, this is not always helpful in making 'binary' decisions such as 'Guilty or Not Guilty?', or 'Influenced by W or Not Influenced by W?'. What engineering (in the above sense) can offer is suggestions of places/patterns of interest that can be investigated 'traditionally' (i.e., narrowing the search space).

But this raises the 'expert-witness' problem again. Even given empirical evidence like '54% of composer X's symphonies contain pattern Y more than Z times in each movement', the *significance* of that fact has to be in some sense a statistical one. In court, the lawyer asks the expert witness: 'Does this, in your professional judgement, mean that X was influenced by W or not?' That process of professional judgement itself involves a statistical process, which has been notoriously misapplied in some recent cases (see, for example, <http://society.guardian.co.uk/nhsperformance/story/0,,1528100,00.html>)

This in turn either suggests that basic statistics should form part of the training of musicologists (as it does for psychologists, for example) or that musicologists should get into the habit of collaborating with statisticians.

A further thread of discussion turned on the fine distinction between 'meaning' and 'explanation' as the goal of music analysis.

e-Science (at the time of the workshop) presents real opportunities in the European academic context. Funding - at a serious level - can be found for the right kinds of project; digital resources (especially in musical audio) are easily accessible; the coming together of applications and ways to tackle the infrastructural problems of coordination are leading to a willingness on the part of politicians to take interest in the potential of Grid computing.

A somewhat prescient view was expressed in a paper by the Bach scholar, Yo Tomita, in 'Breaking the Limits: Some Preliminary Considerations on Introducing an e-Science Model to Source Studies', *Musicology and Globalization*, Proceedings of the International Congress in Shizuoka 2002 (Tokyo: Academia Press, 2004), pp. 233-7.

Aspiration versus reality: Opportunities and challenges in resource discovery

This session was proposed and chaired by Chris Banks (British Library) who began the discussion by outlining certain problems inherent in managing large music libraries that contain both digital and non-digital content. At the British Library, she noted that there was simply insufficient staff to do everything for readers and that therefore systems had to be provided that would allow readers to serve themselves. She

noted the attractiveness of the idea of searching music with music (i.e. content-based music information retrieval). She also pointed out that the British Library holdings are expanding by around 12.5 km of shelf-space per year which places special demands on any systems developed to manage the collections.

As a copyright library, the British Library is obliged to obtain a copy of every published document and this raises a pressing issue with regards to electronic publications: how does the library obtain a copy of every electronic publication without effectively copying the World-Wide Web?

Amanda Glauert observed that it depends partly on whether one considers electronic objects as copies or original documents in their own right. The issue was whether digital documents are 'born digital' or digital transcriptions of non-digital documents.

Matthew Dovey then pointed out various issues relating to the maintenance of electronic publications. In particular, he noted the problems involved in making digital documents persistent (i.e. how to keep them 'alive') and controlling multiple versions of electronic documents. He noted that most electronic documents are in a format that can only be created, edited and viewed using certain software running on certain operating systems. However, all software eventually becomes obsolete which can lead to documents becoming unreadable. One possible solution to this is the use of open-source software. Another is to make documents 'self-describing' - that is, each document defines the format in which it is written.

Michael Casey then suggested that there should be a shift from expecting users to learn how to use library access systems to a culture of expecting library systems to learn how to serve users. Chris Banks noted that this has already started in the form of logs that keep records of searches made on the British Library web site. However, it was clear that there was still a long way to go before music information retrieval systems were truly able to adapt to user requirements. Chris Banks suggested that this might not be possible even in principle as one can never predict how users would want to access the resources. Matthew Dovey suggested that it was a fallacy to believe that a one-size-fits-all design could be found for accessing large music collections that would satisfy everyone. He said that we have to accept that different communities need to access the information in different ways. Raphael Clifford echoed the point by pointing out that there must be thousands of ways in which one might want to search a large-scale music collection. David Meredith suggested the possibility of using machine-learning techniques to cluster users into categories and then automatically generate a different interface for each different type of user.

Raphael Clifford suggested that the data itself could be leased or sold so that different user communities could develop their own interfaces to it.

Geraint Wiggins suggested that a major problem lies in designing appropriate indices and accepting that no single index is adequate to deal with all queries.

Other topics touched on were:

- whether digitization should include analysis of structure of documents or simply provide a 'raw', unparsed image of the resource
- the use of compression for indexing; and
- the problem of some metadata being copyrighted.

Using technology in music psychology research

This session was proposed and chaired by Alexandra Lamont (Keele University), who began the session with a presentation identifying five main areas where technology has and can be used in music psychology research, drawing on examples from recent published and unpublished research from many different settings.

She first considered the use of technology in experimental research contexts. This includes methods used to present musical stimuli that allow for much greater accuracy in information on timing, pitch, timbre and so on, as well as facilitating randomized presentation and automatic trial generation (e.g. E-PRIME software). Technology has also allowed for participant responses to be captured more efficiently and with greater precision (e.g. more precisely measured response times allowing for smaller differences in response times to be considered significant). Technology has also provided completely new methods for gathering psychological data more directly such as various brain-scanning technologies (e.g. fMRI, ERP and MEG).

Lamont then considered the use of technology to build computational models of psychological processes. These models of musical performance, of listener behaviour, and of musical elements and musical structure can allow us to compare predicted behaviour with human behaviour and to attempt to understand learning processes.

Third, technological advances have allowed for improved methods of data analysis in music psychology. Technology enables many different analytic techniques to be applied, including

- analysis of more variable musical inputs such as the human singing voice;
- techniques for measuring and analysing musical parameters such as timbre in real musical performances; and
- video analysis of behavioural performances of and responses to music (for example, using point-light recordings to measure the physical and gestural features of a musical performance or the dancing styles of small children).

Lamont then commented on the influence that web-based technologies have had on the way that music psychology is carried out. In addition to web-based surveys, the Internet can also enable online experimental research with a wider sample. Some research is beginning to explore the nature of online communities in musical composition and in identifying fan groups.

Finally, she observed that research is beginning to capitalize on the technology that people use to play music in everyday life to explore features of everyday engagement (e.g. iPod usage), as well as using technology to capture everyday experiences (random sampling of participants using pagers, mobile phones or PDAs). Video diaries also enable participants to reflect on everyday meaningful experiences without interference from researchers.

There followed a discussion which began with a consideration of the reliability of technological methods and the need for more resources in this area. There was a focus on the rich data available through methods such as video recording. This has had a considerable impact on the kinds of studies that are now possible both in the laboratory and in the field. This brings the advantage of greater possible levels of precision in analysis, as well as the problem(s) of abstraction from the data in order to draw conclusions. The unpredictability of human behaviour was also discussed in relation to modelling approaches, and there is a need for precise modelling approaches that break behaviour down into many different levels which may in turn reduce ethnocentricity. However, although we could argue that all brains are similar, there is also a need to recognize and account for differences between people (both within and between cultures), which comes from the recognition that music is a subjective experience, and the role of the interpreter is a sensitive one.

Expanding the musician's workshop: The effect of information technology on musical creativity

This session was proposed and chaired by Amanda Glauert (Royal Academy of Music). The session began with a presentation from Amanda Glauert which focused on three musicians' 'workshops' at the Royal Academy of Music where technology is intervening in potentially significant ways, and the questions that have arisen or might arise around each one from the viewpoint of creative musicians. It was hoped that by considering these together it would be possible to tease out some common questions that might help move the larger debate forward.

In approaching each of the workshop examples, Glauert suggested that composers and performers would be questioning how new technological tools affected them in terms of:

- the expansion or diminution of their creative materials and creative control;
- the obfuscation or clarification of creative possibilities; and
- the transformation or adaptation of their creative role.

As an institution, the Royal Academy of Music has been investing quite heavily in new technology, both in terms of hardware and specialist staff, some of the latter being themselves high-level composers or performers. The question remains how the institution might bridge the gap between specialists' interests in new technology and the staff/student body as a whole, so as to make new tools a more permanent part of the 'musicians' workshop'.

The first of the workshops considered centred on the topic of electromyography biofeedback for musicians. This technology is one that could be of immediate use to all performers. Melanie Ragge, oboe professor at the Royal Academy of Music and at the Purcell School, has already been collaborating with the physiotherapist Nigel Wilson on pedagogical applications of the new technology. She has found that it is very helpful for oboe students to be able to view their muscle recruitment patterns whilst they are playing, to heighten their awareness of individual strengths and weaknesses, and to lessen undue tensions. The EMG software can produce real-time playback, so that, with a simultaneous audio track, students would be able to hear their performance at the same time as viewing their muscle activity graphically.

While the immediate benefits of such technology to performers might be obvious, the longer-term effects of such new information (and new ways of presenting information) might take longer to evaluate. In particular, one might ask questions such as:

- How might EMG biofeedback affect notions of the relationships between 'sight', 'sound' and 'feel' in performance?
- Would access to such biofeedback information have any interest to composers in releasing new 'inside' information on the physicality of performance?

Such questions of the creative impact of new layers of information become more acute in relation to the second case-study example, the Academy's collaboration with Diana Young from the MIT Media Lab on the development of a 'hyperbow' for cellists. By adding accelerometers, gyroscopes, and force sensors to a conventional bow the 'hyperbow' is able to record the most minute changes in position, acceleration and force applied by the performer. By digitally altering the sound according to such minute changes, Young is offering to expand the possibilities of stringed instruments, and to perfect the control of cellists over their instrument. As a teaching device, the 'hyperbow' can send signals on how the student is holding the bow. As a compositional stimulus, it can be used to enhance sound quality and make it seem as though several cellists are playing at once. According to Patrick Nunn, research-student composer at the Academy, it provides a new way forward for electronic music by, not just adding effects, but allowing them to be controlled by the player's gestures.

The possibilities created by such a hyper-instrument are immediately exciting, but they also raise questions of critical control:

- How could one create a 'repertoire' for the 'hyperbow', that would allow it to be an 'instrument' in the sense of creating a distinctive culture of expectation?
- How might such an instrument be used to enhance the communication chain between composers, performers, and audiences?
- Where are the limits to the 'instrument' that force composers and performers to make choices?

The prospect of a 'blue-skies' world of limitless information was raised in particular by the third case-study example, 'MIDIpedia'. This term has been coined by colleagues at the Academy for the idea of a collaborative site for transcriptions, editions, compositions and arrangements where performers might make their own versions of musical texts or comment on others', and so be able to print out their own tailor-made editions of a huge body of music not currently available in scholarly editions on-line. Such technologically-created means of notating 'versions' or 'editions' of music with the freedom and fluidity hitherto associated with the act of creating 'performances' raises obvious issues:

- What effect might a wikipedia-style approach to editions have to our notion of creative hierarchies?
- How might composers best respond to a MIDIpedia environment?

The common questions raised by these three case studies concern how we might deal with new fluidities in our relation to texts (and notation) and to instruments. Composers and performers recognize that their creative relationships to these are always subject to negotiations with each other, and with audiences. These negotiations require time for reflection, and with the speed of change introduced by new technologies there is not always time to work through what any shifts or changes 'mean'. For a new technological 'toy' to become a 'tool' of creativity requires time for dialogue, both within the creative studio and without.

Representing music on computers

This session was proposed and chaired by Geraint Wiggins (Goldsmiths College, University of London). The discussion revolved around two main questions:

1. What is this 'music' that we aim to represent?
2. What is the 'state of the art' in music representation, and what is generally used?

In this context, we identified a number of needs and desiderata, which will now be discussed.

What is this 'music' that we aim to represent?

This philosophical question is important because it needs to (but does not always) inform practical engineering. Different users mean different things by 'music' and need different things from 'music computer systems'. For example, an audio engineer is likely to mean 'an audio signal' and is likely to need a representation which will facilitate signal processing; a music analyst is likely to need a more structural representation in order to facilitate highlighting intra- and inter-opus relationships.

It is important to identify an appropriate 'musical surface' at which to work for each task and interest: we can perhaps view different needs as being like different 'views' onto some 'musical ideal'.

While these ideas may be thought of as ‘obvious’, there is considerable evidence of a gulf in understanding between those who view music (often in the audio, but sometimes in score or MIDI format) as ‘merely data to be processed’ and those to whom the implicit structure of music is of primary importance.

What is the ‘state of the art’ in music representation, and what is generally used?

There is currently no standard representation system which can maintain the full richness of the artefact being modelled at all levels—that is, all representations are to some extent abstract (e.g. manuscript, audio recording, timbre, texture). Therefore, representations need to allow explicit annotation of structure, for example: declarative representations of music (e.g. performance scores); (multiple) structural descriptions of the ‘data’; (alternative) interpretations of the ‘data’; relationships between, for example, different sources for a given piece; arbitrary (i.e. whatever we need!) meta-data; and linking existing resources together. Above all, representations should facilitate (or at least not obstruct) processing, and they should ideally be usable, at least at the level of the tools that use them, but ideally in themselves also.

Areas currently lacking in consideration are: quality of representation and display, and in particular the physical limits placed on what is possible by transmission and storage media; and broader environmental contextual (meta-)data, such as provenance and media data for non-digital resources.

An important area for the future will be standardization: we would support a ‘music research wikipedia’ to assist sharing of data and representations.

Using technology to analyse musical performance

This session was proposed and chaired by Mark Plumbley (Queen Mary, University of London). The original proposal was to consider the extent to which technology can or could be used to analyse different aspects of a musical performance. Possible issues for discussion included:

- How can digital signal processing (DSP) be used to extract what has been played in a musical performance?
- How reliably could we estimate the identity and timing of notes and chords?
- How, and to what extent, might we measure ‘expressiveness’ in a musical performance?
- What range of ‘sensory modalities’ might we use, including audio, video, MIDI keyboards or special sensors?
- Might we disturb a performance in our attempt to measure and analyse it?
- What applications might there be for such an analysis?
- And perhaps: who might fund research in this area?

Mark Plumbley introduced the session, posing the question of what Digital Signal Processing (DSP) could be used for in analysing musical performance. The discussion that followed converged on several conclusions:

Automatic Music Transcription (i.e. extracting the notes present) would be useful. Limited versions of this are already possible (e.g. transcription of monophonic music or piano) but this is still a difficult problem.

There are many different aspects to a performance. For example, the audience itself may be part of the performance, the movements of the performers, eye contact between performers and audience, etc. might be significant.

It is difficult to know what performance data should be captured - one might ideally want to ‘gather everything’. But there are bound to be ‘blind spots’ in what is measured. Looking back at the data-gathering work of others often indicates what they were interested in measuring. A particular performance only exists once, so one cannot go back to re-record anything missed.

Information about the performance is lost in the recording process. This is not just owing to analogue-to-digital conversion, but also microphone response and placement, acoustics, etc.

Changing aesthetics may change what is regarded as 'expressive'. For example, young students who have grown up with popular music that uses machine-generated, strictly regular rhythms report they prefer strict quantization in timing.

There is a need for a toolkit for music and audio analysis, although it seems unlikely that any single toolkit would do everything required.

Analysis could be used in performance, allowing performers to control the sound or interact in new ways. Ruth Davies gave an example of restoration from old recordings made in 1930s Palestine. Processing was required even before the original recordings could be listened to. The original analogue media (here discs) were fragile, and the digital surrogates produced should last longer than the originals. However, long-term preservation of digital data requires on-going active maintenance.

Precision tools, fuzzy concepts, and ill-formed questions

This session was proposed and chaired by Alan Marsden (Lancaster University) who began with an introductory talk in which he laid out the problem of resolving the apparent incompatibility between the precision of computational procedures and the imprecision of the questions, concepts and methods of musicology. He suggested that musicological questions are often ill-formed and fuzzy, such as, for example, the question 'What key did Mozart associate with death?'. He proposed that, when faced with such a question, we could start by turning this imprecise question into properly empirical hypotheses, such as:

- Hypothesis A: Death words imply X minor.
- Hypothesis B: X minor implies death words.

Marsden observed that, to test hypothesis A, one would require a tool that can find the perceived key at any given point in a musical passage; whereas, to test hypothesis B, one would need a tool that can reliably track the perceived sense of key throughout a musical passage. Marsden pointed out that, currently, methods exist for carrying out the first of these tasks considerably more accurately than the second. However, most currently available tools for key analysis do not model reliably and accurately the musicologists' concept of key. A similar problem exists for many musicological concepts that are not systematically defined and for which computer tools can therefore only provide approximately accurate models.

Given these difficulties, Marsden proposed three possible ways to proceed. The first approach would be to try to make software tools that model expert musicological thinking much more accurately. The results generated by such tools would then have some authority and the process of developing them would teach us a great deal about expert music cognition. However, developing such accurate computational models is time-consuming and it may be practically impossible to develop tools that model expert music cognition sufficiently accurately for them to be useful.

The second possible strategy would be to use computers to do well-defined tasks and then rely on experts to infer interesting generalizations and conclusions from the results generated by these well-defined processes. For example, one could use a computer to produce a trace of the pitch-class frequency distribution within a moving window throughout a piece and then infer from this various conclusions about the perceived key at each point in the music. This places the responsibility for determining the local key back in the hands of the expert musicologist but provides the expert with extra data to better inform his or her decision.

The third possible strategy would be to embrace the fuzziness and uncertainty in musicological ideas and theories by using fuzzy or statistical computational approaches instead of more discrete, categorical models. Such methods would be more likely to mimic closely the actual paradigms of current musicology by accepting the reality of fuzzy concepts and ill-formed questions.

There followed a discussion which began with the observation that empirical disciplines (e.g. experimental psychology) have well-established techniques for measuring error and taking it into account. However, it is often not immediately obvious what an 'error' is in musicology as there is rarely a firm ground truth against which one's predictions can be evaluated.

Nevertheless, musicologists routinely deal with error (e.g. copyists' errors in manuscript sources, misprints in many nineteenth-century printed editions), so it might not be too hard for musicologists to incorporate more well-developed error-handling techniques into their methodologies. Also, developers of software tools for musicology should be explicit about the degree of error to be expected in a tool's output (e.g. the degree of uncertainty in the predictions made by a key-finding algorithm).

In a similar way, those who compile digital collections of musical sources, such as encodings of scores, should estimate and quote the degree of error introduced in the encoding process.

It was observed that computers are also often used as case-discovery tools rather than analysis tools. For example, it would be natural to attempt to use a computational technique to find cases where Mozart associates a specified key with death. Software tools for information retrieval can be extremely useful even if they do not exhibit perfect recall and precision (consider, for example, the widespread use of the Google search engine).

Another important technique to consider when analysing musical data is that of using several different tools for the same job and combining their results using, for example, a simple voting system in order to obtain a more accurate result than could have been obtained from any of the tools individually. If different approaches produce compatible results, then the problems of inaccurate tools, fuzzy concepts and ill-formed questions are cancelled out to some degree.

This last point highlights the importance of knowing precisely what a given tool aims to do and how it does it.

Musicological arguments typically use various different paradigms and methodologies and have various different purposes and goals. Such arguments almost invariably involve some degree of interpretation but the goal might be to recognize truth or to recognize importance or quality; and the research might aim to highlight particular phenomena rather than find general rules.

Either way, when using computer tools, it is particularly important to be aware of one's assumptions since the machine will process the information uncritically. Therefore, perhaps musicologists should embrace a culture more like that adopted in empirical sciences, of admitting ignorance and pursuing clarity - a culture in which one gives the best available answer to a problem and acknowledges that a better answer might be possible in the future.

Summary and conclusions

The principal aim of this workshop was to allow experts in music with an interest in technology to talk to experts in technology with an interest in music and identify ways in which they can collaborate fruitfully to achieve worthwhile goals. Many of the participants seem to have found the day useful for making new acquaintances, exploring possibilities for future collaborations and finding out what sort of work is being done in other fields.

One recurring issue was the mismatch between the methods used in science and those used in the humanities. The perception was that scientific methods involve precise methods but that humanities scholarship involves fuzzy answers based on ambiguous evidence. It emerged that one way of reconciling the two approaches might be for musicologists to start using statistical methods in a more sophisticated way. It was also suggested that statistics might become an essential subject of study for those interested in pursuing musicological research.

There is now a huge amount of musical information in digital form and the challenge is to make this information accessible to users in appropriate ways. Unfortunately, different communities need to interact with the same musical information in different ways, which suggests that the best strategy might be to allow users to create or customize their own interfaces on the data (cf. Google's customized home page facility). Another issue with respect to digital information is that of guaranteeing its persistence, preserving its accessibility and availability in the face of changing operating systems, file-formats and software and communication standards.

In music psychology, the use of ICT has allowed researchers to have much more control over their experiments, measure features of responses much more precisely and store and analyse the data generated much more easily. Computer-modelling of perceptual processes has also allowed complex psychological theories to be tested rigorously.

In conservatoires such as the Royal Academy, composers and performers have been experimenting extensively with the use of ICT to enhance and transcend traditional musical practice. But it seems that some time is required for practitioners to evaluate and explore the potential of new technologies before we know whether or not they will be ultimately useful and valuable.

The way that musical information is represented in computers critically limits the ways that it can be used. It is therefore important to identify an appropriate 'musical surface' (e.g. digital audio, notation, MIDI, harmonic analysis, etc.) for each application. Each different musical surface gives a different 'view' onto a particular 'musical ideal'. An important area for future research will be the standardization of musical formats at different structural levels and the reliable conversion between these formats.

There are various ways in which digital signal processing (DSP) techniques can fruitfully be used in the analysis and control of musical performance. There is clearly a need for a usable toolkit for music and audio analysis. However, such a toolkit would have to overcome the problems inherent in the fact that musical surfaces exist at various different levels of structure (e.g. audio, notation, MIDI, etc.). Also, when using recordings of a musical event for research purposes in, for example, psychology or ethnomusicology, one must often deal with issues such as non-optimal acoustic conditions and losses of information in the recording process.

As mentioned above, there seems to be an incompatibility between the precision of the data provided by computational methods and the ambiguity and provisional nature of musicological 'hypotheses'. Musicologists could overcome this by recasting vague hypotheses as more precise, empirically testable ones. However, many central musicological and music-theoretical concepts do not have universally-agreed definitions (e.g. 'key', 'cadence', 'harmonic structure') which means that there are often many different ways of translating vague musicological claims into testable scientific hypotheses. It also makes it difficult to design tools for automatic extraction of musical structure that all musicologists would find useful since any such tool would have to implement particular precise explications of ambiguously-expressed theories. A promising alternative, therefore, might be to use computers to carry out well-defined tasks such as statistical analyses and then use the results generated to support or refute more generally-expressed musical theories. A third possibility is to use new techniques such as machine-learning and Bayesian approaches to attempt to model directly the fuzziness in musicological theory.



Acknowledgements

The author would like to thank Tim Crawford, Amanda Glauert, Alexandra Lamont, Alan Marsden, Mark Plumbley, Frans Wiering and Geraint Wiggins for providing notes and text for the sections in this document that summarize the discussion sessions.