

Towards the compleat musicologist?

Prof. Nicholas Cook, FBA

Professorial Research Fellow, Royal Holloway, University of London & Director, AHRC Research Centre for the History and Analysis of Recorded Music
[Invited talk at ISMIR 2005]

I

We stand at a moment of opportunity. It's an easy thing to say. We stood at a moment of opportunity around 1980, in terms of the relationship between music theory and psychology: each discipline had good reasons for being interested in the other—for psychologists because music is a complex, culturally embedded activity that is open to quantitative analysis, for music theorists because empirical methods allowed them to ground and develop work that until then had been basically speculative. The opportunity was realized and the result was a gain for both partners. A new subdiscipline developed, but strongly linked to its parent disciplines and with new insights and ways of working that fed back into both.

Obviously I wouldn't be saying this if I didn't think there was a similar opportunity today for musicologists in relation to music information science. I'm choosing my words carefully because the interdisciplinary interface I'm talking about here is more complicated than the music theory/psychology one of the 1980s. On the music side I'm talking about musicology in the broad sense that includes theory, ethnomusicology, and the rest. On the information science side I'm talking about a range of empirical and computational work that draws on and feeds back into computer science, cognitive science, acoustic engineering and I don't know how many other subdisciplines. I don't know how far there can be said to be such a thing as a discrete discipline, or subdiscipline, of music information retrieval, but for the musicologist wanting to expand his or her ways of working it seems a more complex and to that extent daunting field than was the case of music theory/psychology twenty years ago.

I hope that that music information scientists feel that there is an opportunity for them in a closer relationship with musicology; I'm sure that a closer relationship with music information science presents musicologists with an opportunity. The trouble is, we've been standing at this moment of opportunity for quite some time now. By the end of the last century a great deal of work had gone into developing a variety—a huge variety—of music-computational tools, and quite a number of studies making use of these tools had been published. Naturally enough, such studies, and I'm thinking for instance of the work of David Huron and his co-workers, were designed to illustrate the potential of the tools: that's to say their starting point was what the tools could do rather than a musicological problem crying out for solution. What we still don't have to any considerable extent is the mirror image, that is to say studies that are grounded in mainstream musicological problems and that make use of computational tools as simply one of the ways you do musicology. In fact the area of musicology from which you would most expect such a response, what on the continent is called systematic musicology, has only a marginal existence within the Anglo-American discipline; such empirical musicology as there is tends to go under the strange and misleading label 'music psychology'. The ball, in short, is firmly in the musicologists' court, and there's a danger that the opportunity will be missed.

In this paper I shall discuss what I think are some of the factors inhibiting musicologists' engagement with information science—though this means I'll also be talking about what I think are some of the opportunities, because the inhibiting factors and the opportunities are

really different sides of the same coins. I'm going to be talking about ways of thinking, ways of construing music as information, ways of using the information you get out of it—the point being that we're *all* engaged in music information retrieval, but within very different frameworks, and the better we all understand that the better the chances for meaningful interaction will be. (Even in the case of the more straightforward relationship between music theory and psychology there were lots of opportunities for misunderstanding, because people were talking about the same things but for different purposes and on different assumptions.)

However I'm going to start briefly with something that happily *isn't* the problem: hostility to technology on the part of musicologists. The way in which audio-video presentation technology has developed has seen to that: it simply doesn't make sense to teach music without technology nowadays, so that virtually everyone who teaches music either has mastered the technology or feels guilty that they haven't. And in the last few years so many crucial sorts of musical information have come on line that without really knowing it musicologists have become used to doing more and more of their research from their desktops: I'm talking about established resources like RISM or RILM or J-STOR, as well as newer or developing resources like the CHARM (that's AHRC Research Centre for the History and Analysis of Recorded Music) on-line discography, DIAMM (Digital Archive of Medieval Music) and OCVE (On-line Chopin Variorum Edition). The last three of these are delivering primary musicological materials to your desktop—though CHARM's ability to deliver streamed audio is threatened by proposed copyright legislation on historical recordings—while OCVE goes further in that it enables you to manipulate passages from Chopin sources on screen, in effect replicating what you might do on paper. And some of us also use utilities which embody complex symbolical processing such as Themefinder, an electronic thematic dictionary that is actually running Humdrum, but protects the user from it since it's a utility designed for a single purpose. So the idea of doing musicological research at your desktop is well established.

But there's a difference between changes in working practices and changes in the nature of what you do, and all this established usage stops just short of what I'd call computational musicology proper—which I see as involving sufficient understanding of the symbolical processing and data representation on which it's based that you're not locked into a single pre-defined purpose, but can develop or combine your ways of working in line with your particular research purposes. As things stand at present, it's a bit like going from driving a car to knowing what happens under the bonnet, and if that's really the sticking point (and I'll be coming back to that) then it would be very easy to say that what's needed is new training opportunities—which of course is true, and I'll be coming back to that as well. But I also think it's necessary to appreciate some underlying conceptual factors—factors that have to do with how musicologists work, and that can be traced back to the historical development of the discipline—so it is to these that I now turn. And I shall illustrate what I say by reference to a project that I'm just now starting within CHARM, in conjunction with Craig Sapp: this involves working with large numbers of recordings of Chopin's Mazurkas, a distinctive repertory of 58 dance-based piano pieces which have been frequently recorded since the days of wax discs and cylinders, and it represents just one possible take on the opportunities that I've referred to.

II

Basically I'm going to make two points that have to do with data poverty, as well as saying something about what might rather grandly be called epistemological issues. The first point is that *musicologists are used to working with highly reduced data*. Mainly, of course, I mean scores, which are such drastically simplified representations of musical sound that you almost

want to say they *symbolize* rather than *represent* it: people don't play musical rhythms as written, often they don't play the pitches as written, and that's not because they play it wrong but because that the notation is only an approximation. And that's before we start thinking about all those dimensions like timbre and texture that aren't directly represented in the notation at all. All of these missing elements have to be supplied by the performer or the musicologist if you're to make sense of the score as music. In the absence of that there's a real sense in which you're studying scores and not music, and there's also a real sense in which that's what traditional musicology was set up to do: it was part of the nineteenth-century nationalist project, and early musicologists saw themselves as musical philologists. Their job was to reconstruct their national repertoires by editing and interpreting early musical sources, and so they thought of music as in essence of form of writing, abstracted from its contexts of performance and consumption. Or you might say that they thought of it as a kind of literature, with musical 'works' being conceived on very much the same basis as literary works. In fact you could say that they had a very narrow view of what music information was.

Now when I say that musicology is oriented towards the study of scores I don't of course mean that musicologists don't care about how music sounds. But whereas the score is tangible the sound is intangible: the data you can actually manipulate is highly reduced. And this makes for problems when you start manipulating scores as if they were mathematical formulae rather than cultural objects dependant on contextual interpretation, for example in the case of set theory (an approach to the analysis of twentieth-century 'art' music developed by Allen Forte which—to put it bluntly but I think not inaccurately—means translating scores into numbers and then doing maths with them). Even when musicologists work directly with musical sound, as in analysing recorded performances—which has recently been quite a growth area—there is a strong tendency to reduce the data in such a way as effectively to end up with a new kind of score, and I'll go into this in a little more detail.

One of the simplest but most widespread ways in which computers have been used in recent musicology is through the use of tempo graphs, most often made by what is generally called the 'tapping' method, though I prefer to think of it as 'reverse conducting': you listen to the music and tap on a computer keyboard at regular intervals, for instance once a bar, with the computer logging the times and then turning them into a graph. There are quite a number of problems with this technique, which I won't go into, but the basic point I'm trying to make is how very reduced the data is as a representation of what's actually going in the performance. So how does the 'Mazurkas' project compare? You could say it's a half-way house. There have recently been a lot of developments in onset detection for piano music, and we'll be using a semi-automated system developed by Andrew Earis which gives you timing and dynamic information down to single-note level; we'll be correlating the timing and dynamic information with the scores, using established pattern-matching approaches. In other words we'll be looking for motivic, harmonic, or structural patterns both within individual mazurkas and across the repertory as a whole, and using these as a basis for analysing the timing and dynamic information. And we'll be interested in things like how far we find regularities within individual mazurkas, across groups of mazurkas, or the repertory as a whole, or again how far they correlate with date of recording, or the pianist's date of birth, or pianistic 'schools', or individual performers.

This is, as I said, a half-way house, in the sense that one could instead be trying to exploit the new audio-based approaches being developed for purposes of genre classification and playlist generation. Actually we have plans to do just in collaboration with Mark Sandler's team here at Queen Mary, University of London. But there's another sense in which it's not a halfway house at all, because the use of the musical score as a basis for analysis has a particular significance for musicology, and at this point I'm going to stray briefly into epistemology.

When you listen to music you don't have an internalized score that you match the sounds to, and that's why the inspirational project on piano performance directed at OFAI by Gerhard Widmer hasn't use scores in the way I've described: instead the approach has been a bottom-up one giving rise on the one hand to very low-level stylistic rules, and on the other to extremely high-level characterisations of different great pianists' styles. Literate musical cultures, however, like that of Western 'art' music, involve a constant tension between what is heard on the one hand and what is constructed or manipulated through notation on the other: as a technology of the imagination, music writing stretches people's powers of perception, and the history of that stretching is in a sense the history of Western 'art' music. So for the musicologist the score *isn't* just a highly reduced and therefore unsatisfactory representation of music: it's also a fundamental aspect of the culture that gives rise to the music, and in this way the emphasis on notation in the 'Mazurkas' project isn't a kind of stopgap but an essential part of what it means to study the music *musicologically*. There is in this sense a complementarity between the 'Mazurkas' and the OMFAI approaches: they're working towards different ends and within different epistemological frameworks—or to put it another way, they're construing music as different sorts of information.

To get back to the point, musicologists are used to working with highly reduced data, but I've illustrated through the example of analysing recordings how there's a movement towards engagement with full-sound data, and this is obviously going to be more and more important as we increasingly deal with music in the age of web-based multimedia—music in which conventional notation rarely plays the same cultural role that it used to. It will always be the business of musicology to understand sound in terms of the means of representation that define the culture, but I suspect we shall increasingly need to draw our basic ontological categories from outside the traditional discipline. However that may be, the conclusion is clear: *working with fuller data will upon up new areas of musicology*.

But that's only one of the two aspects of data poverty that I wanted to talk about, the second of which is that *musicologists are used to working with small data sets*. Basically what I'm talking about here is the issue of comparison: you might see the principle of comparison as fundamental to all musical analysis, and yet musicologists have for generations been in a state of denial about it. A hundred years ago comparison lay at the heart of musicology. Style was seen as a fundamental musicological issue, and it was studied through the comparison of musics from different times and places, with the intention of extrapolating fundamental musical principles that applied everywhere and distinguishing them from those variable features that defined the music of different nationalities or historical periods. But this whole project got entangled with extreme right-wing ideologies, for instance with attempts to define the essential qualities of Aryan tone-consciousness, and after the Second World War it became seen as a thoroughly imperialist exercise, comparable to the ethnographic museums in London or Paris where relics artefacts from the farthest reaches of the empire were collected, compared, and assigned to different stages in the evolution of man (*sic*) or civilization. 'Comparative musicology' was replaced by ethnomusicology, the basic principle of which is that all cultural productions have to be understood in the context of their own society and that to make comparisons across cultures is illegitimate. In the same way though rather later, 'style analysis' was replaced by structural analysis, according to which a particular music pattern means nothing in itself, but only in its particular structural context—from which it follows that it is illegitimate to make comparisons between different pieces, each of which has to be considered on its own terms.

Although I'm again putting it rather crudely, this is the basic context within which musicologists have thought it proper to work very closely—some might say myopically—with very small data sets, usually a single musical work, or even a movement. Perversely, this

means that computers, which made it feasible to compare large bodies of data, came along just as that kind of approach was going off the musicological agenda: that's why the earliest examples of computational musicology, such as Arthur Mendel's work of the 1960s on Josquin attributions, are at the same time some of the last examples of traditional style analysis. So the idea in the 'Mazurkas' project of working with large numbers of recordings across a complete repertory, and so focussing on issues of compositional and performance style, in some ways represents an attempt to reinvent an older musicological approach, an older construal of music as information.

But this is something that's in any case being driven by the move towards analysing performances to which I referred. The kind of performance analysis to which I referred, often involving the correlation of a single performance with the structural features visible in the score, in effect assumes that each performance is created directly out of the score as if for the first time. But that's now how performers perform or listeners listen. For performers and listeners, each performance exists not only 'vertically', so to speak, in relation to the score, but also 'horizontally', in relation to other performances; to put it another way, performances are relational. So the 'Mazurkas' project is analysing each recorded performance of a particular work, or of the repertory as a whole, against all the others. One way to make the point is this: the conventional tempo graph measures the performance in relation to a metronomic norm, in relation to the kind of mechanical performance a sequencer can create but a human can't. But when you put it that way, it's not at all obvious why this is a sensible baseline against which to view the individual performance. It might make more sense to create a grand average tempo profile from all available recordings of a given mazurka and compare individual recordings to that. Or perhaps more productively, you could develop a generative model of the music—for instance by combining the mazurka 'script' (with its elongated second beat) with Neil Todd's phrase-based model and Johann Sundberg's more surface-oriented model, as implemented in his programme 'Director Musices': by adjusting the relative values of these three inputs and perhaps their internal settings as well, you'd try to match the individual performance you were studying, resulting in a baseline against which its genuinely idiosyncratic features would show up in much greater detail than in the conventional tempo graph. And at the same time the settings you'd used in this would become a means of representing the relationships between different performance styles.

However this works out in practice, I think the conclusion is again clear: *working with larger data sets will open up new areas of musicology*. But at this point I'd like to bring up another epistemological issue. There may be a few musicologists who are actively hostile to computational approaches, but there are many more who simply don't see the point, and I think a major reason for this is the feeling that musicology is based on the *experience* of music—and that once you bring computers in the experience goes out, meaning that the whole exercise becomes pointless. But there's a basic confusion between technology and epistemology here. What I've described in the 'Mazurkas' project is not a computational substitute for the experience-based musicology I referred to, and I'll make the point through two illustrations.

First, as a cultural musicologist I'm interested in the way in which the Mazurkas acquired particular connotations of gender and exoticism in the nineteenth and early twentieth centuries, and I want to understand what particular elements of their style were associated with these values, or how performance style interacted with compositional style in the construction and maintenance of these associations: the better an understanding of style I have the better I will be able to do this, and that's why I need methods for analysing large data sets. (Otherwise anything I say about style is likely to be selective, speculative, or plain wrong.) As for the second illustration, this goes back to what I was saying about creating

baselines against which to identify the idiosyncratic features of a particular recorded performance. That's not using the computer as a substitute for experiencing the music: on the contrary, it's a means by which you become able to see and, more to the point, *hear* the music with greater precision and sensitivity, to enhance your experience of it. Using methods derived from information science, then, doesn't mean giving up on the traditional epistemological values of musicology, on the sort of knowledge you hope to acquire from it. Working with fuller data and larger data sets can open up new areas of musicology, but it can also mean doing traditional musicology better. I'm not sure that, without a working knowledge of such approaches, you can really claim to be the compleat musicologist today. I certainly hope you won't be able to tomorrow.

Twenty years ago there was a tendency to think of music-analytical approaches as dogmas: to be a Schenkerian analyst you had not only to believe in Schenkerian theory but also to believe that all other approaches were false. My generation of theorists ditched that idea in favour of seeing Schenkerian analysis as a large, bulky tool, a sort of power drill if you like, which can be used in a lot of situations but not when you need a saw or a hammer. I think we're in the same situation with computational musicology that we were with analysis twenty years ago, and that musicologists aren't going to be comfortable using computational methods until they have come to see them the same way, as elements of the musicological toolkit. In fact, I think it might be a good idea if we stopped talking about 'computational musicology' at all, and instead just talked about doing musicology with computers.

III

So how might all this be translated into action? In this paper I've been emphasizing the conceptual changes that will be necessary if computational approaches are really to take off within the musicology as a whole, though of course you could equally say that it's the spread of computational methods that will promote the conceptual changes, so there's something of a chicken-and-egg situation here. I'm sure that things that will need to happen on both the information-technological and the musicological sides if the opportunities that I spoke about at the beginning of this paper are to be realized.

One good thing is that more and more music is becoming available on the web, whether in score, MIDI, or audio form, along with more and more utilities for format conversion—if, that is, you know where to find them (and even in the age of Google that can be a big 'if'). From the musicologist's point of view I think the main thing that makes the prospect daunting is the variety of quite separately conceived computational tools, interfaces, and music representation languages, all of which adds up to a learning curve that is hard to accommodate with all the other demands of an academic career. Probably what I shall call the jobbing musicologist's dream would be one single system that did everything, like the 1990s turnkey office solution, but of course that isn't going to happen, not only because there's no Microsoft to fund it, but also because of the sheer variety of completely different things that people want to do with musical data. All the same, current developments in linking and annotating different representations of music, and new search technologies for audio and video materials, could help to integrate computational approaches into musicological working practices in a way that is not the case at present; to the extent that these developments are taken up commercially, they may change the way in which musicologists do things almost without our realizing it, in the same way that developments in multimedia computing have revolutionized the way we work on film music or music video. Equally important is the potential for bringing together different music-computational tools within a unified interface, as in the OMRAS2 project proposed by Mark Sandler and Tim Crawford.

But what will be critical from the jobbing musicologist's point of view is the trickling down of research, so to speak—the translation of cutting-edge research into practical, usable, everyday tools. It's hard for example to persuade specialists in music computing of the value of graphic user interfaces. I know that if you're using Humdrum every day then it's much quicker to use the raw Unix interface, but jobbing musicologists are going to be using Humdrum in April and then not again until August, and for that kind of user even something as simple as Michael Taylor's 'Humgui' makes a huge difference, because it reminds you what the different commands are and how their syntax works (and even what they're called). A more sophisticated interface like Andreas Kornstädt's 'J-Ring' is potentially much better, of course, because the on-screen notation taps into musicologists' established skills and ways of working.

Even the fanciest GUI, however, isn't likely to relieve us from knowing what's going on under the bonnet—for example, in the case of Humdrum, how the music is represented and therefore how the different available operations can be invoked to achieve a desired result. So I can see no way round the learning curve, in other words the issue of training, and probably what is going to be critical here is what happens in terms of new recruits to the discipline. I don't see how computational approaches are going to be assimilated into the musicological mainstream until basic concepts and skills are embedded in training programmes taken by all musicology postgraduates, not just music computing specialists: I'm talking about basic principles of music representation, how to operationalize musicological problems, how to use different computational tools in conjunction with one another. Having said all that, however, I still have doubts. Is this necessarily the right approach? Instead of training musicologists, should we be promoting patterns of collaboration between musicologists and music computing specialists? But then, how could the musicologists work effectively with the music computing specialists unless they already knew how to operationalize musicological problems? And how would such collaborations be funded? How would the music computing specialists develop their career routes? How would we avoid the very hierarchical relationships that currently characterize humanities research carried out by a 'principal investigator' in conjunction with 'research assistants'?

I don't know the answer to these questions, and I'm not aware of their being talked about, even here in London, which—if you put together the different institutions where computational musicology of one sort or another is going on—must rank as one of the world centres for such work. All the potential for a major disciplinary advance in musicology is there, but we've got to put in place the conditions to make it actually happen. Otherwise we could be standing at this moment of opportunity for a long time to come.