



City Research Online

City St George's, University of London

Citation: Cottrell, S.J. (2019). Big Music Data, Musicology, and the Study of Recorded Music: Three Case Studies. *Musical Quarterly*, 101(2-3), pp. 216-243. doi: 10.1093/musqtl/gdy013

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/19984/>

Link to published version: <https://doi.org/10.1093/musqtl/gdy013>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Big Music Data, Musicology, and the Study of Recorded Music: Three Case Studies

Abstract

This paper considers some of the interactions between Music Information Retrieval (MIR) and musicology, particularly in relation to Big Music Data and the analysis of recorded music. Since MIR is still not widely recognized within the musicological community, and the possible insights offered by analyzing Big Music Data even less so, the paper both briefly contextualizes some of this work for a musicological readership and provides three specific case studies that illustrate concrete musicological outcomes. These relate to: changing orchestral pitch over time; pulse salience beyond the Euro-American classical music tradition; and changing performance tempi in classical music. The paper concludes by considering some broader conceptual issues that arise from the relationship between Big Music Data, musicology and recorded music.

Author Biography

Stephen Cottrell is Professor of Music at City, University of London. His research interests encompass three inter-related areas: ethnographic approaches to musicians and music-making, particularly within the Western classical music tradition; the study of musical instruments, especially the saxophone; and the study of musical performance. His publications include *Professional Music-making in London* (2004) and *The Saxophone* (2013). During a freelance career spanning nearly two decades he earned an international reputation as a saxophonist performing contemporary music, particularly as leader of the Delta Saxophone Quartet.

Key Words

Big Music Data, Musicology, Recorded Music, Digital Music Lab, Music Information Retrieval

Introduction

Music Information Retrieval (MIR)—a term that embraces a range of interdisciplinary but science-based approaches to gain information about music—remains for many musicologists a rather enigmatic part of the music studies field.¹ Some of the reasons for this are considered below. The more recent arrival of Big Music Data (BMD)—very large-scale collections of data pertaining to music in the form of audio recordings, consumer preferences, bibliographic information, etc.—is an emergent topic that, as David Huron argues, contributes to the transformation of musicology from a data-poor to a data-rich field.²

This paper considers the relationships between MIR, BMD and musicology, particularly in relation to the analysis of recorded music. It starts by providing a brief overview of the study of recorded music in the different parts of the music studies field, before giving some insights into ongoing MIR work and how this supplements and augments musicological work elsewhere. The main part of the paper considers the opportunities afforded by software such as that generated as part of the recently-established Digital Music Lab, as well as reflecting on both the technical and conceptual challenges that BMD provide in relation to musicology. Three case studies are detailed, which evidence the type of musicological insights that can be gained from investigating BMD. The paper concludes by introducing Franco Moretti's concept of 'distant reading' as a hermeneutic device to explain the possible benefits of BMD for musicologists.

Music Studies, Recorded Music and Music Information Retrieval

The academic study of recorded music and its artefacts—which I have elsewhere characterized as 'phonomusicology'³—may be characterized as falling into three broad but related categories. The first considers the cultural and social contexts in which recorded music is produced.⁴ The second focuses on

the recorded musical sound itself, and seeks, through the application of different analytical strategies, to infer information about performance practice and/or the theoretical principles underpinning the music system captured in the recording.⁵ The third comprises the materials and conduits through which recordings are disseminated.⁶ Categories two and three are particularly relevant to this paper, although it is the analytical ‘mining’ of recorded musical sound that is most significant here.

This analysis of recorded music has been inflected differently within the various parts of the music studies field. It is perhaps longest established in the sub-field of ethnomusicology, which from the late nineteenth century until the 1950s was known as comparative musicology and was largely founded upon the transcription and analysis of recordings of ‘non-western’ music. The act of transcription was particularly prioritized during this period, as evidenced by Bruno Nettl’s observation that ‘the ethnomusicologist for long was in the first instance a transcriber of music. The first task of the field was thought by some to be the transcription of all available recordings.’⁷ This early relationship between Comparative Musicology and recorded music is set out in more detail elsewhere,⁸ but the significant points here are that evolving audio technology was driving the field; that one of the benefits of engaging with this technology was seen to be the manner in which it more easily facilitated the production of transcriptions (a word that circulates with slightly different inflections throughout both the MIR and musicology communities); and that direct comparisons between different musics were at the heart of the intellectual endeavor. One of the subtexts of this paper is that not only did this kind of comparative musicology never really disappear, but it is also increasingly resurgent. One final connection with ethnomusicology’s historical roots is that comparative musicologists saw themselves as akin to laboratory-based scientists and were seeking to emulate the kinds of detailed empirical work being undertaken in the natural sciences. This is not to suggest that any part of the music studies field is today seeking to return to nineteenth-century traditions of grand, overarching scientific narratives. It is just that methodologies and mind-sets of a particular type of contemporary science—in this case

computer science—are similarly exerting an increasing influence on one corner of the musicological landscape.

Other parts of the field demonstrate rather different relationships with recorded music. Those studying popular music or jazz have little option but to engage with recorded sound. Just as comparative musicologists were usually concerned with musical cultures that had little or no use of music notation, so too are the aural and contextual characteristics of popular musics often seen as being of greater significance than any literary or musical texts underpinning them. In contrast, musicologists concerned with Western classical music have only more recently embraced the opportunities offered through studying recordings, building on Robert Philip's ground-breaking research,⁹ and augmented through work such as that undertaken by the UK's CHARM project.¹⁰ It is the close analysis of music in performance which recorded sound facilitates that has significantly underpinned the 'performative turn' in recent musicology which, for example, Nicholas Cook identifies.¹¹

These thumbnail sketches demonstrate that the computational analysis of recorded music now cuts across sub-disciplinary boundaries. It might be argued that the roots of this unificatory trend lie with the melograph, an automatic transcription machine that has its roots in the nineteenth century, but which was particularly advocated by Charles Seeger in the 1950s. The details pertaining to this are well covered elsewhere,¹² but we might note Seeger's prescient assertion in 1958 that 'the automatic graph can serve as a bridge between musics—a common denominator, as it were'.¹³ By extension, therefore, the automatic graph could serve as a bridge between musicologies also. Naturally, Seeger could not have foreseen how technological developments would unfold, but with computational methodologies now being put to the service of a range of music analyses involving a variety of music cultures, this bridge between musicologies is precisely what we are witnessing.¹⁴

Given the insights that MIR potentially offers, it is perhaps surprising that these computationally-driven approaches have not been more widely adopted within 'mainstream'

musicology. Instead they remain favored only by a small core of musicologists who are familiar with the field and the tools available, and by a slightly larger number of scholars who would perhaps consider themselves computer scientists as much as musicologists. Yet MIR studies often replicate and augment the kinds of humanistic questions with which most musicologists are frequently concerned. Recent examples—offered here only to underline this point for those unfamiliar with such work—might include genre classification in Latin music based on rhythmic pattern matching; assessments of jazz compositional style based on identification of chord progressions; or identifying classical music genres through automatic detection of key progressions.¹⁵ Although MIR work has often focused on Western classical and popular styles (in part because computational musicologists themselves are more familiar with these styles and the software is better able to deal with equally tempered scales), more work has recently been undertaken on global folk and traditional musics, particularly with respect to India and Turkey.¹⁶ And while the current paper is concerned with computational analyses of audio data, there is also extensive work being undertaken on the analysis of other forms of musical representation, such as musical scores, for example in the computational analysis of dissonance in the music of Palestrina and Victoria.¹⁷

With so much research now being undertaken in MIR it is reasonable to ask why the outcomes arising from it are not having more impact on musicology. Inevitably, the reasons are complex. Perhaps first among these is that musicology, as a discipline, has developed over time particular ways of engaging with and thinking about music; these approaches are then replicated by those neophytes entering the field who become inculcated in its ways.¹⁸ Whether such approaches are philologically inflected studies of musical texts or ethnographically grounded investigations of music and context, the empirically driven methodologies favored by the hard sciences, and particularly the more data oriented insights of computational science, are not necessarily easily integrated into long-established musicological worldviews.¹⁹ There is perhaps a residual musicological mindset that ‘values’ or

‘meanings’ remain secreted in musical scores, related primary sources, or collective acts of music-making, in ways that only the human brain can properly comprehend. Furthermore, because of their mathematical bases, MIR publications also tend to be replete with unfamiliar acronyms and puzzling mathematical formulae, and as Alan Marsden observes, ‘musicians, unfortunately, are rarely equipped to understand a mathematical formulation.’²⁰ The interfaces for some MIR software are not always especially user-friendly for the uninitiated, when compared with the more familiar programs encountered on personal computers, supported as the latter are by multi-million-pound development budgets of which MIR researchers can only dream.²¹ MIR software also tends to be bespoke, designed to solve a particular MIR query, and is not usually intended to be adapted to address other questions. Sometimes, too, MIR research is more focused on solving computational challenges, for which music notation or musical sound provide useful case studies, to push the boundaries of computational science rather than musicological knowledge; this is a perfectly laudable aim, but not one with which musicologists are necessarily empathetic or equipped to support.²²

But there are perhaps other historical dynamics at work also. The reinvigoration of musicology that occurred during the 1980s and ‘90s, through the work of scholars such as Joseph Kerman, Susan McClary, Lydia Goehr or Marcia Citron, recontextualized the position of music theory within the discipline at large, positioning such approaches among those drawing on critical and cultural theory in a manner in which the analysis of specific patterns of musical notes did not always play a central role.²³ While close readings of the musical score remained—and continue to remain—important, the previous centrality given to these approaches was arguably diluted by the introduction of other critical perspectives. Yet it was over the same approximate timescale that computational approaches to music analysis began to develop, greatly extending the potential empirical detail and depth of such analyses. It might therefore be argued that the different groups of scholars were on slightly different trajectories: the detailed empirical insights offered by MIR methodologies, and the analytical insights they often

provide, were being made available just at a time when many musicologists, operating in a now more fragmented musicological landscape, were perhaps not quite so interested in what they had to offer.

Big Music Data—musicological issues and challenges

The application of MIR technology to Big Music Data (BMD) represents a significant opportunity for both computer scientists and musicologists. The vast quantities of music audio now stored in digital repositories, and the potential insights offered through the examination of these in the ‘super-human’ manner that computers afford, arguably represents a paradigm shift for musicology. Through these large collections of audio data, questions about human music-making can be interrogated on a scale not previously conceived.

But what is Big Data and how might it be defined in relation to recorded music? Mark A. Beyer and Douglas Laney suggest that Big Data comprises ‘high volume, high velocity, and/or high variety information assets that require new forms of processing to enable enhanced decision making, insight discovery and process optimization’.²⁴ One of these ‘new forms of processing’ is considered below, but the question of scale is complex in relation to recorded music. On the one hand, while popular music collections may contain a few million recordings, and classical collections tens of thousands, ethnographic collections are often small, arising perhaps from a few dozen and occasionally a few hundred recordings. These numbers might not constitute Big Data as it would be understood in other contexts. On the other hand, any single audio track comprises millions of data points. Every recording has a considerable number of pitch events, for example, as well as rhythmic and timbral data. Even a collection of a few hundred recordings contains a very large quantity of data, to the extent that it might be described as ‘big’ and interrogated accordingly. So the question of what constitutes Big Data in

relation to recorded music is one of perspective, and related both to the units of data under consideration and the kinds of questions one might wish to ask.

Outside of academia, the major drivers of MIR research in relation to BMD have been allied to popular music, for several reasons, not least of which is the personal tastes for pop music among researchers leading the field and the relatively large datasets to which they therefore had access. Furthermore, Western studio-oriented pop music tends to be heavily diatonic and is often produced—with respect to its tonal clarity and/or harmonic simplicity—in ways that make it inherently sympathetic to the kinds of interrogation that MIR software facilitates. And the commercial imperatives of the music industries, combined with the potential profits to be made by harnessing MIR-generated insights into popular music consumption, have meant that greater resources have been allocated to support development in this area. Familiar apps such as Shazam or Soundhound, and music recommendation algorithms embedded in sites such as iTunes or Amazon—‘if you liked that then you might like this’—are now widely used to identify music or track musical tastes, largely in the hope of persuading users to purchase material related to or resembling music for which they have already demonstrated a preference.²⁵

Yet it would be wrong to think that the activities of these companies are entirely commercially oriented. The music streaming company Spotify, in particular, recognizes that it can mine its download data for information about taste preferences and consumption patterns, and has created an MIR section within the company—its Insights team—which undertakes work that is essentially musicological.²⁶ For example, on 24 December 2015 the company added the Beatles back catalogue to its streaming service. 100 days later it disseminated an analysis of the downloads of 250 million Beatles tracks by age, gender, and geographic location, showing which tunes were most popular in which categories. The infographic provided on the company’s website facilitated zooming in to city and even street level, clearly evidencing the global consumption patterns for this material over this period (and

demonstrating that, ironically, Liverpool—the home of the Beatles—was one of the UK cities slowest to access it).²⁷ While musicologists may be concerned about the validity of these results—because they have no opportunity to access the raw data and thus cannot verify the analysis—they might also be envious of the sheer scale of data that the company was able to draw on, and the speed with which their corporate resources were able to analyze it and disseminate the results.

Although this particular example is quite musicological in nature, many of the algorithmically-driven strategies of these online corporations—the new majors, as Ben Mandelson has described them²⁸—are rather different, since they are essentially ad hoc exercises in musical futurology, which try to predict and shape individual and collective taste preferences for the purpose of eliciting greater capital returns. As Jeremy Wade Morris points out, these providers can be seen—from a Bourdieuvian perspective—as information-dependent cultural intermediaries (he uses the term ‘infomediaries’), who endow the products whose circulation they control with varying amounts of symbolic and cultural capitals. They are not ‘neutral purveyors of predictions’ since they ‘exert a logistical power that shapes the ways audiences discover, use and experience cultural content’.²⁹ Nor are these entirely automated processes, contrary to appearances, since they rely on the interaction of company staff with the algorithms to differentiate and distinguish the approach of one company from another.³⁰ How these individuals monitor and manipulate musical taste would undoubtedly provide an interesting musicological study, but their interventionist activities are qualitatively different from the post hoc analyses offered here and elsewhere, in which the traces of human music-making, sedimented in musical scores, recordings and so forth, are considered for the evidence they provide of previous musical practices.

Because computational analysis requires digitized sources, the current extent and availability of digital repositories present particular challenges. Although Western popular and classical music genres are now extensively digitized, they are also covered by stringent intellectual property rights (IPR)

regulations that may make access to them difficult for researchers. Like other Big Data repositories, they are often in private hands rather than the public sector, and thus controlled by corporations whose interests do not necessarily align with the ostensibly altruistic investigations of the curious scholar.³¹ And although many recent recordings are digitized, older recordings may not be, and thus comparisons of changes over time (for example) are more difficult. Music traditions beyond these genres are inevitably less well represented. Many archives around the world still have most of their materials in analogue form. Digitization is time consuming and expensive, especially of the order of magnitude that might then be described as constituting Big Music Data. Archives such as those at the British Library or the Smithsonian Institute are digitizing and making available some of their collections, but these are also enmeshed within a range of IPR considerations. Accessing both the files and the metadata can be problematic, and although overall global numbers are significant, bringing them together in an accessible fashion is problematic.

Another challenge lies with the metadata itself. Every track in data sets such as these is usually accompanied by information describing the content: the names of the performers or the music, perhaps where and when it was recorded, etc. This information—the metadata—is often inconsistently inputted. The same word may be spelled differently, some information may be omitted or put into different fields, the genre may be wrongly assigned, etc. This metadata can be cleaned up manually by a researcher working only with a small number of tracks, and it can still be done, albeit laboriously, for a few hundred. But making tens of thousands of metadata listings consistent is an altogether different proposition.³²

Ethnographic collections are often characterized by recordings that have both speech and music, either because the musical performance—frequently recorded in the field—is framed by speech on the part of the participants, or because the ethnographer was also conducting interviews, or has verbally

recorded information about the event. Speech segments such as these obviously risk impacting on the computational analyses.

Finally, it should be borne in mind that extracting meaningful data from audio is generally difficult, even in contexts where much is known about the music and/or its attendant organizing principles and performance aesthetics. Without this, the obstacles become yet greater. There remains not only a degree of technological fallibility in the manner in which computers infer information about musical characteristics such as beat intervals, tempo relationships or pitch levels, but also a recognized disjunction between the ways in which computers ‘hear’ music and how the human brain processes it. Some of these issues will be returned to in the case studies below.

But perhaps the greater challenges are intellectual rather than technological. What meaningful questions can be asked of such large music collections when they encompass so many diverse styles and/or periods and places? How do we get beyond a disciplinary mindset in which music-making is so frequently associated with named individuals (composers, performers) or specific localities, yet still provide useful insights into musical creativity?

Notwithstanding these challenges, the musicological mining of audio Big Music Data has occasionally been undertaken by others, if seldom by those who might describe themselves primarily as musicologists. For example, both Serra et al. and Mauch et al. have sought to demonstrate what they see as evolutionary trends in popular music, the former employing methodologies drawn from statistical physics and complex networks, the latter using approaches found in evolutionary biology.³³ Perhaps because of their intellectual heritage, both papers, while very rigorously worked out, rely extensively on the kinds of complex mathematical formulae that, as already noted, humanities scholars sometimes find off-putting.³⁴

Nevertheless, the following section provides three more case studies illustrating some musicological insights that might be achieved through the examination of audio Big Music Data, and further attempts at addressing some of the challenges it presents.

Musicology, Big Music Data and the Digital Music Lab: Three Case Studies

The Digital Music Lab (DML) was established to develop the research methodologies and appropriate software required to analyze large-scale music collections.³⁵ The system presently facilitates access to four data sets, of different sizes and dispositions. The largest is the commercial audio collection known as *I Like Music*. This has some 1.6 million recordings, covering many different genres but largely comprising popular and library music; around 288,000 tracks are presently available for analysis by the DML. The digitized components of the British Library collection comprise some 51,000 tracks, covering largely Western classical music and ethnographic recordings. Two smaller collections—4,800 classical recordings digitized as part of the CHARM project³⁶ and a small collection of 2,700 Chopin piano recordings from the Mazurka project (itself a sub-set of CHARM)³⁷—are also available. In total, therefore, there are some 346,000 tracks available for analysis.

The DML system has servers embedded at the British Library and *I Like Music*. IPR restrictions mean that in many cases the recordings cannot be taken off site. In such cases the DML system exports only the results of the analysis and not the audio itself, which remains located on the host server. In some cases, however, the audio is free of IPR restrictions, and can be streamed through the DML system on request, thus enabling the user to focus on a particular track in the analysis.

While the DML gives access an embarrassment of audio riches, as already noted, it presents challenges in endeavoring to ask musicologically meaningful questions across such a large and

disparate corpus of recordings. The three case studies outlined below demonstrate further ways in which those challenges might be addressed.

Case Study 1. Changing pitch levels over time

The pitch levels used in classical music performance have varied widely over time and place. Information about such pitch levels in the centuries prior to the invention of recordings must be inferred from measuring legacy artefacts such as instruments or tuning forks, or from written evidence asserting the use of a specific pitch in a particular context; all these primary sources have their drawbacks.³⁸ With recordings, however, it becomes possible to measure empirically the pitch used for a given performance. But this is an area fraught with difficulty, particularly when relying on earlier forms of sound-reproducing technology in which manual crank-handles would be used to drive rotation, the speed of which directly determined the pitch levels recorded and heard. Even after the introduction of electrical recording processes and electrified machine rotation in the mid-1920s, pitch levels could still be inconsistent. But this has not prevented some scholars from endeavoring to use these early recordings as evidence. In his consideration of orchestral pitch levels in the twentieth century, for example, Bruce Haynes lists ‘a number of recordings made in Germany between 1920 and 1943’³⁹ which, he argues, show considerable discrepancies in the pitch levels found in each recording. But relying on individual recordings in this way is risky, since the technological limitations of these systems at the times the recordings were made, combined with potential inconsistencies in the transfer process when making modern copies of the original recordings, can all undermine the validity of tuning data measured from any one recording.

Here we see both the potential for and current limitations of a Big Music Data approach. In theory, if one had a very large number of 78 rpm transfers from which to work, individual

idiosyncrasies in relation to the transfer of any one recording might be ameliorated by the very large number of measurements that could otherwise be made; the data set would be of sufficient magnitude that trends over time should in any case emerge, and discrepancies and inconsistencies might reasonably be seen as outliers to those overall trends (providing, of course, that one has broad confidence in the manner in which the data set has been generated). But the lack of commercial imperative for the digitization of recordings from the 78-rpm era (with a few notable exceptions for well-known artists such as Caruso), and the seemingly prohibitive costs of doing so from the perspective of publicly-funded institutions, means that the numbers presently available for analysis remain rather small, and we are not yet at the stage where we can aggregate the results from tens of thousands of recordings, as would be desirable.


Nevertheless, it is possible to demonstrate certain trends empirically, and to consider whether, as is often asserted, there has been a general rise in orchestral pitch over the nineteenth and twentieth centuries. This trend was widely discussed in the nineteenth century,⁴⁰ but Haynes notes that as late as the 1990s ‘there remained a general impression among musicians that pitch was rising’.⁴¹ To what extent might this be verified using BMD analysis?

Some intriguing possibilities arise from the DML. Table 1 shows a comparison of pitch levels for recordings by orchestras and quartets (overwhelmingly string quartets) over the course of the twentieth century.

Table 1

		1900-10	1911-20	1921-30	1931-40	1941-50	1951-60	1961-70	1971-80	1981-90	1991-2000
Symphony Orchestra	No. of recordings	7	6	109	103	180	260	30	67	77	467
	mean pitch	443.0	441.2	443.3	444.1	444.2	444.0	444.9	443.9	445.1	443.7
	SD	4.11	0.86	2.84	2.92	2.53	2.50	1.34	1.50	1.55	2.47
Quartet	No. of recordings	34 ⁴²	42	76	86	94	36	44	56	46	233
	mean pitch	443.8	441.8	442.4	443.7	443.4	444.3	443.2	444.1	444.3	444.4
	SD	1.9	3.65	2.49	1.74	2.15	0.84	1.3	1.28	1.52	1.46

Because there are relatively few recordings from the first two decades of the century, and bearing in mind the technological limitations of the recording equipment and the variabilities already noted, it is perhaps worth setting these aside and considering only the period from 1921-2000. Plotting these data on a scatter graph, as shown in Figure 1, suggests that pitch levels may indeed have risen over the century.

Figure 1

Caption: Pitch levels of symphony orchestras and string quartets 1921-2000

The quartet line is particularly consistent from 1960-2000 and implies an average increase in pitch of approximately 1 Hz over this period (just under 4 cents). The point is significant because string quartets are unconstrained by the physical capacities of the human voice (singers will often complain that the constant raising of pitch levels strains their voices)⁴³ nor are they ‘anchored’ by the pitch levels customarily associated with the piano; the tuning is determined entirely by the string players. Furthermore, over this timescale, recording fidelity—now captured very largely on LPs and CDs—is sufficiently accurate for us to have considerable confidence in these figures. The orchestral figures are more variable but can still be seen to demonstrate increasing pitch levels over the course of the century. The notably lower figure registered for the 1990s, and the larger standard deviation, may reflect the influence of historically-informed performances at this time, since such groups often play at lower pitch levels than prevailing symphonic norms.

Why overall pitch has seemingly risen in this fashion is inevitably a matter for conjecture. It may reflect the premium that has increasingly been put on brighter, edgier sounds in recent decades, something that was already in train at the beginning of the century, when steel strings began to be preferred by many string players to the gut strings that had previously been the norm. It is often averred that string players tune slightly above wind and brass players, in order that the string sound appears more brilliant, and this may have driven overall pitch levels higher. There may have been some transfer of performance aesthetics from various popular music genres, especially those involving amplification. It may also have been a consequence of the improvement in recording fidelity in the second half of the century, with the crispness and clarity increasingly found on recordings driving changing timbral aspirations from performers, again particularly for string players. Indeed, the impact of technologically-

mediated listening practices is likely to have played a significant role in overall expectations of appropriate classical music sounds.

Case study 2. Cross-cultural pulse salience

Even within the somewhat circumscribed recorded music datasets that are presently available for analysis, there are some tantalizing insights into the possibilities of cross-cultural music comparisons. As noted above, earlier forms of comparative musicology were similarly driven by evolving audio technology, and it is worth recalling that gaining cross-cultural understandings about human musicality from analyzing phonograph recordings was already underpinning the work of scholars such as Carl Stumpf in the late nineteenth and early twentieth centuries. Stumpf used such recordings to investigate principles of melodic and rhythmic construction, tonal variety and instruments etc., among different global music cultures, as part of his psychologically-inflected investigations into *The Origins of Music* (*Die Anfänge der Musik*).⁴⁴ But such investigations are potentially contentious, since the broader the geographical sweep of recordings analyzed, the closer one inevitably comes to discussions of human universals in music-making. Such discussions have proved controversial among scholars, particularly ethnomusicologists, who have historically cherished the idea that particular patterns of musical sound are most meaningful for the groups or communities that produce them, and have argued that this cultural specificity makes it difficult (or impossible) to generalize about human music-making on a global scale.⁴⁵ Cognitive psychologists, however, tend to take a different view, and have argued for the biological bases of shared human musical behavior. Dale Harwood, for example, asserts that such music universals as might be identified are ‘not to be found in specific musical structure or function [but] are examples of basic human cognitive and social processes at work’.⁴⁶


One of the basic cognitive processes to have received attention involves the way human beings understand musical pulse, and particularly how the brain segments pulse information into specific metric groupings.⁴⁷ Arising from this has been the suggestion that there are optimal underlying pulse speeds at which we find it easiest to identify prevailing metric structures. Richard Parncutt, for example, suggests that ‘pulse sensations are most salient in the vicinity of 100 events per minute (moderate tempo)’.⁴⁸ Justin London similarly asserts that ‘the special significance of periodicities in the 600-700 millisecond range [i.e. 100-85 beats per minute] has long been known’.⁴⁹

Nevertheless, a recurring feature of music psychology experiments is that the participants in such studies tend to be recruited from those Euro-American contexts where the studies are frequently undertaken. In Parncutt’s oft-cited study, for example, the 22 participants were drawn from students at the University of Stockholm and researchers at Stockholm’s Royal Institute of Technology. Perusing back issues of a journal such as *Psychology of Music* demonstrates the predominance of Euro-American contexts for music psychology studies. It also reveals the centrality of particular Western instruments such as the piano in these studies, not only because such resources are close at hand but also because, in the case of the piano, the instrument readily lends itself to the detailed measurements that are necessary for the kinds of studies of music perception and cognition that are at the heart of this sub-field. None of this is to denigrate any of the important work being undertaken in this area, it is simply to suggest that claims about basic cognitive processes might be usefully illuminated by the opportunity afforded by Big Music Data to interrogate recorded sound across different music cultures. If the human brain is at its most comfortable processing metric patterns from underlying pulses of around 100 beats per minute, it might reasonably be assumed that music-making cross-culturally would show a greater preponderance of patterns using pulses around this speed. It would make sense that the underlying pulses adopted in any given music culture, while undoubtedly socially negotiated, would

also have a biological basis: human beings are more likely to be drawn to things that the brain finds easier to process than those that are cognitively more challenging.

Computers can help identify musical beats in several ways. The simplest approach is for a programme to record taps made by somebody listening in real time, so that those tapping patterns can then be analyzed. More sophisticated software will interrogate the audio files directly, determining automatically where the musical beats appear to fall. The bespoke nature of the first approach would be prohibitively labor intensive for large collections, and the DML software contains an automatic beat tracker which assesses the large-scale distribution of beats per minute (BPM) in any given collection of recordings. The software maps the frequency of these BPM, showing how many occur within any given range, and illustrates this using a tempo histogram. Figure 2 demonstrates this for the British Library collection using the keyword ‘Ireland’. This brings up 622 recordings of Irish music, which the software then analyses collectively.

Figure 2




Caption: Tempo histogram for British Library ‘Ireland’ audio collection–622 recordings

The highest peak here shows that there are more than 14,000 instances of BPM occurring in the range of 114.1-119.4 (the exact figure is 14,664). The other bars show lower values for different BPM ranges. Each category can then be seen as a percentage of the total number of BPM identified, thus demonstrating their relative weighting within the collection as a whole. In this collection, 153,057 beat events were identified in total, and thus the most frequently encountered category (114.1-119.4) accounts for just under 10% of that total.

Taking the same approach to other collections enables us to make comparisons between them. Figure 3 demonstrates this for seven audio collections at the British Library, varying between 77 recordings of Indian music to the 1000 recordings of music from Suffolk, England.⁵⁰ In each case the BPM have been mapped as percentages of the overall number identified.

Figure 3



Caption: BPM distribution for seven British Library audio collections (figures in brackets indicate the number of recordings analyzed in each collection).⁵¹

As is evident from the graph, there are clusters of most frequently encountered BPMs in a range from about 104 up to around 137 or so, before in all cases tailing off. This potentially casts a slightly different light on the assertions made by Parncutt and London. While the base mark of ‘around 100 beats per minute’ would seem congruent with the lower end of the range implied by these findings, this analysis suggests that other music cultures feel equally comfortable processing music patterns with average underlying pulse speeds above this; and, while being cautious about the risk of cultural stereotyping, this appears yet more notable in those cultures where drum patterns play particularly important musical roles, such as India and Ghana.

Nevertheless, there are some important caveats to note here. First, while the DML software can rapidly identify what it thinks is the underlying pulse across a large number of audio files, computational pulse perception may be rather different from human pulse perception. And as will become apparent in the next case study, the beat tracking element of the DML software tends towards identifying faster pulses than appears reasonable; to put it another way, the computer identifies smaller intervals between what it thinks are ‘beats’ than the human brain would. These differences could only be properly identified under laboratory conditions, but in relation to BMD this would require laboratory work on an unfeasible scale. Second, these figures show only the relative prevalence of what are assumed to be the underlying pulses in each collection. They demonstrate nothing about how metrical groupings might be perceived by human beings in these culture groups, since that would also require detailed ethnographic work.

Such analytic approaches to Big Music Data are emergent forms of scholarship and are offered here at the level of suggestion rather than assertion. At the very least, this analysis suggests that the

range of periodicities with which the human brain feels most comfortable is greater than that suggested by Parncutt and others, and that more work needs to be done to understand the relationship between musical patterns most commonly found in different cultural areas and their impact on the development of the brain. But the point here is not to support or refute categorically the views noted above. It is rather to demonstrate the potential of BMD to inform cross-culturally our understanding of human music-making, such that we might respond to Steven Brown and Joseph Jordania's plea for 'a new movement in comparative musicology to take full advantage of the amazing musical database generated by ethnomusicologists',⁵² while recognizing that at present there is still some way to go before such insights can be made with confidence.

*Case Study 3. Tempo changes in classical music over the 20th century.*⁵³

One of the more substantial shifts registered in musicology over the past quarter of a century has been an increasing focus on the study of musical performance, and in particular a move away from studying classical music *and* performance to one that involves studying classical music *as* performance.⁵⁴ The analysis of recorded music has played a central role in these developments, for much the same reason that comparative musicology of the late nineteenth and early twentieth century was also heavily reliant on recorded music: recordings can be played repeatedly and thus the performance characteristics they preserve can be analyzed more closely; live performances are, by their very nature, more transient and thus more difficult to analyze in detail (unless they are themselves recorded).


The examination of performance tempi constitutes a recurrent theme through this body of scholarship, and it is frequently asserted that performance tempi have, in general, slowed down as the century has progressed. Jose Bowen, for example, notes that after World War II, performances of the second movement of Beethoven's Fifth Symphony have tended to play the repetition of the second

theme (heard loudly in C major and dominated by brass instruments) slower than its original hearing (heard softly in the tonic Ab major and dominated by the woodwind).⁵⁵ Bowen sees this as evidence of a tendency after the War increasingly to associate Beethoven and his music with ideas of victory and heroism, whereas previously he had been seen as ‘pompous’.⁵⁶ While there may be specific reasons for this shift, Robert Philip sees the slowing of performance tempi as being widespread, noting that ‘maximum tempos within movements are usually slower in post-war than in pre-war performances, so that the average tempo of a movement has generally dropped. And what is true within movements is true also of complete multi-movement works’.⁵⁷ As he points out, part of the reason for this may be attributed to the reduced recording time available on 78 rpm records from the early decades of the century, and the consequent need for performers to speed up their tempi to accommodate the music on one side of a record. After 1925, with the advent of electrical recording, and particularly from the LP era onwards, such technological constraints were gradually obviated.

If there has been a general reduction in performance speeds in Western classical music performance over the course of the 20th century, one would expect to see this reflected in BMD analyses of performance tempi, but the picture offered here is more mixed. Figure 4 shows the average tempi calculated from performances of works by four composers—Mozart, Beethoven, Schubert, and Brahms—from 1960 to 2000. As noted in case study 1, the date range chosen here avoids the recording and reproduction inconsistencies inherent in 78 rpm records and earlier technologies, the fluctuations in which risk having a particular impact on the calculation of performance tempi; this date range covers recordings that would have been released on LP or CD. And as noted in case study 2, the absolute values recorded here appear improbably high in relation to the tempos we might perceive as listeners; the computer is ‘hearing’ the music differently from the human brain. While this would be important if one were making claims about any given performance or work, it is less important for making

comparisons on this scale: such ‘irregularities’ can be assumed to be broadly consistent over the many hundreds of recordings under scrutiny.

Figure 4



Caption: average tempi per year for performances of music by Mozart, Beethoven, Schubert and Brahms, 1960-2000.⁵⁸

The trendlines suggest that performance tempi in works by Beethoven and Brahms have decreased over the period, in keeping with observations noted by Bowen and Phillip. The increasing tempi in Mozart performances appear to contradict these assertions, however, and may well demonstrate the greater impact of historically-informed interpretations, in which performance speeds are often notably faster than other approaches. But the similar apparent increase in performances of

works by Schubert is worthy of note. Schubert and Beethoven are often seen by musicologists as having both classical and romantic qualities in their music.⁵⁹ Nevertheless, such views are usually informed by analytical readings of the musical texts; hence Nicholas Cook's call for renewed emphasis on the study of music *as* performance, which may reveal different insights from those based on textual analysis alone. Perhaps these figures suggest that Schubert has been aligned by *performers* more frequently with classical composers, whereas Beethoven's works have more often been seen as fundamentally romantic and interpreted accordingly.

Conclusions

The study of Big Music Data contributes to the ongoing reinvigoration and expansion of the music studies field, and particularly the development of more empirical approaches to musicology.⁶⁰ The case studies outlined here demonstrate both the potential insights that a comparative musicology of BMD might bring in relation to the study of recorded music, as well as illustrating some of the current challenges and limitations of these approaches. Such challenges are not unique to the study of recordings. In their analysis of an equally large corpus of musical-bibliographic data, Stephen Rose, Sandra Tuppen and Loukia Drosopoulou focus especially on the changing nature of paratextual material—such as title pages—in mid-sixteenth century music anthologies.⁶¹ They note similar challenges in relation to metadata alignment and cleaning, intellectual property rights, and access restrictions. While their study is obviously rather different from the one outlined here, both evidence a belief that, as they put it, 'the analysis of large datasets allows new ways of studying music history'.⁶²

They also draw attention to Franco Moretti's concept of 'distant reading', in which he distinguishes between the inferences that may be drawn from examining large corpora of material and those that arise from the 'close reading' of individual texts.⁶³ It is the latter that more usually

characterize humanistic study, including musicology. Close reading encourages the creation of canons, because it relies on the examination of a corpus of texts produced by those deemed to be the most important figures of whatever literary, artistic or musical tradition is under scrutiny. The history of musicology is, by and large, a history of canon construction. As Moretti points out, ‘if you want to look beyond the canon [...] close reading will not do it. It’s not designed to do it, it’s designed to do the opposite’.⁶⁴ In making assertions of the large-scale nature of changing cultural patterns through the examination of BMD, whether pertaining to the paratextual material of sixteenth-century music anthologies or the rising pitch of twentieth-century classical music performance, the perspectives necessarily required to draw meaningful conclusions may dilute the importance of a particular subset of notable works or composers, however legitimate the arguments might be in other contexts for their canonicity.⁶⁵

This is an emergent field, and speculation on future directions naturally risks being swiftly outdated. Nevertheless, BMD offer significant possibilities for our understanding of human music-making on a global scale. The geo-tagging of audio files—i.e. having metadata that tells us more precisely where recordings come from—may allow us to map musical activity internationally (as Spotify has already been able to do for those areas connected by the Internet to its repository); from this we might be able to demonstrate empirically, for example, the global distribution of particular kinds of scales—e.g. pentatonic or heptatonic scales—even when, as will normally be the case, they are not exactly the same scales. As further music repertoires become digitized we may more easily be able to observe the change of specific harmonic or rhythmic patterns in particular corpora of music. We may also be able to better identify large-scale changes in performance practice over time, as opposed to changes in compositional practice, building on the more bespoke work already undertaken by several scholars noted above.

Notwithstanding these intriguing possibilities, it would be wrong to imply that the increasing digitization of the world's music represents a musicological utopia in which we will soon be able to know anything about everything. The difficulties of accessing and engaging with these collections have already been noted, and there are profound moral and political implications in the way these musical resources are accrued, stored and controlled by those 'new majors' who own the rights to them. As Andrejevic et al. point out, the concatenation of these huge data sets risks inscribing 'a new digital divide between those with access to data (and the resources for making sense of it) and those without'.⁶⁶ Digital mega-corporations such as Google, Facebook and Amazon 'can experiment constantly for commercial purposes on their (sometimes very large) user populations without notifying them, whereas those supposed to be operating in the public interest are effectively locked out or provided only limited access.' Ethnomusicologists will be particularly familiar with such concerns, echoing as they do the discourses of both anxiety and celebration that accompanied the explosion of commercial interest in World Music in the 1980s,⁶⁷ and the potential disregard for individual and collective intellectual property rights on the part of the global music corporations that was perceived to accompany it.⁶⁸

Following Moretti, it might be argued that the analysis of BMD offers the potential to circumvent traditional disciplinary focuses either on individuals whose outputs are often seen as canonical (whether Beethoven, Brahms, Madonna, or Björk, for example), or music cultures whose work is elevated through the work of particular scholars (e.g. the Kaluli people of Papua New Guinea or the Venda of Southern Africa). But the way BMD are being utilized by some global corporations, particularly in relation to popular music, is already leading to the creation of alternative canons. Recommendation algorithms, while purportedly responding to individual user preferences in a manner that may be seen as quintessentially democratic, arguably drive users towards a very small proportion of specific artists or tracks, generating not only canons that can then be rapidly capitalized upon by the

companies themselves, but also a ‘long tail’ of musical material that risks being marginalized.⁶⁹ Ideally, the perspective offered here, and the size of the collections under consideration, is sufficiently inclusive that even largely unacknowledged musical material can be included in the data set and thus the analysis.⁷⁰

As already noted, the computational and comparative approaches to music analysis that inevitably inhere in engaging with BMD are resonant with the comparative musicology of earlier times. Such approaches will not appeal to everybody, in part because they tend to de-emphasize some of the key concerns on which modern musicology is predicated. But as Moretti observes in relation to comparative literature, ‘there is no other justification for the study of world literature [...] but this: to be a thorn in the side, a permanent intellectual challenge to national literatures—especially the local literature. If comparative literature is not this, it’s nothing.’⁷¹ Engaging with Big Music Data offers the possibility of a comparative musicology played out on a super-human scale. It may yet prove to be similarly discomfiting in the challenges it offers to certain long-cherished musicological tenets.

¹ I am deeply grateful to Emmanouil Benetos, Tim Crawford, Christophe Rhodes and Tillman Weyde for their comments on a draft version of this paper, and for the insights offered by an anonymous reviewer.

² See David Huron, ‘The New Empiricism: Systematic Musicology in a Postmodern Age,’ in *The 1999 Ernest Bloch Lectures* (2001), 16.

³ Stephen Cottrell, ‘The Rise and Rise of Phonomusicology,’ in *Recorded Music: Performance, Culture and Technology*, ed. Amanda Bayley (Cambridge, UK: Cambridge University Press, 2010), 15-36.

⁴ For example, Amanda Bayley, *Recorded Music: Performance, Culture and Technology* (Cambridge, UK, 2010); Louise Meintjes, *Sound of Africa!: Making Music Zulu in a South African Studio* (Durham, 2003); Jonathan Sterne, *The Audible Past: Cultural Origins of Sound Reproduction* (Durham, 2003); Jonathan Sterne, *Mp3: The Meaning of a Format* (Durham, 2012).

⁵ For example, José A. Bowen, 'Finding the Music in Musicology: Performance History and Musical Works,' in *Rethinking Music*, ed. Nicholas Cook and Mark Everist (Oxford: Oxford University Press, 1999), 424-51; Daniel Leech-Wilkinson, 'Using Recordings to Study Musical Performance,' in *Aural History*, ed. Andy Linehan (London: The British Library, 2001), 1-12; 'Portamento and Musical Meaning,' *Journal of Musicological Research* 25 (2006), 233-61; Robert Witmer, 'Stability in Blackfoot Songs, 1909-1968,' in *Ethnomusicology and Modern Music History*, ed. Stephen Blum, Philip V. Bohlman, and Daniel M. Neuman (Urbana: University of Illinois Press, 1993), 242-53.

⁶ For example, Nicholas Cook, 'The Domestic Gesamtkunstwerk, or Record Sleeves and Reception,' in *Composition, Performance, Reception: Studies in the Creative Process in Music*, ed. Wyndham Thomas (Aldershot: Ashgate Publishing Limited, 1998), 105-17; Kyle Devine, 'Decomposed: A Political Ecology of Music,' *Popular Music* 34, (2015), 367-89; Jonathan Sterne, *Mp3: The Meaning of a Format* (Durham, 2012).

⁷ Bruno Nettl, *The Study of Ethnomusicology: 29 Issues and Concepts* (Urbana, 1983), 67.

⁸ John Baily, 'Modi Operandi in the Making of 'World Music' Recordings,' in *Recorded Music: Society, Technology and Performance*, ed. Amanda Bayley (Cambridge: Cambridge University Press, 2010), 107-24; Kay Kaufman Shelemay, 'Recording Technology, the Record Industry, and Ethnomusicological Scholarship,' in *Comparative Musicology and the Anthropology of Music*, ed. Bruno Nettl and Philip V. Bohlman (Chicago: University of Chicago Press, 1991), 277-92.

⁹ Robert Philip, *Early Recordings and Musical Style: Changing Tastes in Instrumental Performance* (Cambridge, 1992).

¹⁰ The AHRC Centre for the History and Analysis of Recorded Music was a five-year research project established on 1 April 2004. Full details are available at <http://www.charm.rhul.ac.uk/index.html> (accessed 7 November 2017).

¹¹ Nicholas Cook, *Beyond the Score: Music as Performance* (Oxford, 2015).

¹² Peter Crossley-Holland, 'Special Issue on the Melograph,' *Selected Reports in Ethnomusicology* 2/1, (1974), 2-120.

¹³ Charles Seeger, 'Prescriptive and Descriptive Music-Writing,' *Musical Quarterly* 44, (1958), 184-95, at 195.

¹⁴ See also Nicholas Cook, 'Bridging the Unbridgeable? Empirical Musicology and Interdisciplinary Performance Studies,' in *Taking It to the Bridge: Music as Performance*, ed. Richard Pettengill and Nicholas Cook (Ann Arbor: University of Michigan Press, 2013), 70-84.

¹⁵ Tlacaél Miguel Esparza, Juan Pablo Bello, and Eric J. Humphrey, 'From Genre Classification to Rhythm Similarity: Computational and Musicological Insights,' *Journal of New Music Research* 44, (2015), 39-57. Thomas Hedges, Pierre Roy, and François Pachet, 'Predicting the Composer and Style of Jazz Chord Progressions'. *ibid.* 43, (2014), 276-90. Christof Weiß and Maximilian Schaab, 'On the Impact of Key Detection Performance for Identifying Classical Music Styles' (paper presented at the 16th ISMIR Conference, Malaga, Spain, 2015). Many more examples can be found in *The Journal of New Music Research* or the annual proceedings of the *International Society of Music Information Retrieval*.

¹⁶ This particular regional focus arises largely from the Compmusic project directed by Xavier Serra in Barcelona: <http://compmusic.upf.edu/> (accessed 7 November 2017). See also the special issue on 'Computational Approaches to the Art Music Traditions of India and Turkey, *Journal of New Music Research*, Vol 43/1, 2014.

¹⁷ Andie Sigler, Jon Wild, and Eliot Handelman, 'Schematizing the Treatment of Dissonance in 16th-Century Counterpoint' (paper presented at the 16th ISMIR Conference, Malaga, Spain, 2015). It might be noted that there are many other programs that offer similar functionality in relation to interrogating symbolic music data (rather than audio files), such as Michael Cuthbert's open access *music21* programme (<http://web.mit.edu/music21/doc/index.html>)

¹⁸ See also Pierre Bourdieu, *Homo Academicus*, trans. Peter Collier (Cambridge, 1988), 84-90.

¹⁹ There are of course other parts of the music studies field—notably music psychology and in computational approaches to music analysis—to which this sentence does not apply.

²⁰ Alan Marsden, 'Response to Guerino Mazzola,' *Journal of Mathematics and Music* 6, (2012), 103-06, at 104.

²¹ See also Nicholas Cook, 'Computational and Comparative Musicology,' in *Empirical Musicology: Aims, Methods, Prospects*, ed. Eric Clarke and Nicholas Cook (Oxford: Oxford University Press, 2004), 107.

²² For further discussion on some of these issues see also http://transforming-musicology.org/blog/2015-11-27_addressing-the-music-information-needs-of-musicologists/

²³ Joseph Kerman, *Contemplating Music: Challenges to Musicology* (Cambridge, MA, 1985); Susan McClary, *Feminine Endings: Music, Gender, and Sexuality* (Minnesota, 1991); Lydia Goehr, *The Imaginary Museum of Musical Works: An Essay in the Philosophy of Music* (Oxford, 1992); Marcia J. Citron, *Gender and the Musical Canon* (Cambridge England ; New York, 1993).

²⁴ M.A. Beyer and D. Laney, 'The Importance of Big Data: A Definition,' (Stamford, CT, 2012).

²⁵ These commercial applications work in different ways. Shazam uses audio recognition, analyzing the musical sound to determine which recording it is from and using this to identify music that the software believes to have some sonic similarity. Retailers such as iTunes or Amazon are analyzing consumer preferences allied to catalogue metadata. As with any Big Data project, the quality of that metadata influences the success of any searches based on it. See Gregory Camp, 'Spotify,' *Journal of the Society for American Music* 9, (2015), 375-78.

²⁶ <https://insights.spotify.com/uk/>

²⁷ See <https://insights.spotify.com/uk/2016/01/22/meet-the-beatles-data/> (accessed 7 November 2017).

²⁸ Stephen Cottrell, 'An Interview with Ben Mandelson,' *Ethnomusicology Forum* 19, (2010), 57-68, at 67.

²⁹ Jeremy Wade Morris, 'Curation by Code: Infomediaries and the Data Mining of Taste,' *European Journal of Cultural Studies* 18, (2015), 446–63, at 447.

³⁰ *Ibid.*, 456.

³¹ See also Mark Andrejevic, Alison Hearn, and Helen Kennedy, 'Cultural Studies of Data Mining,' *ibid.*, 379-94, at 380.

³² See also Stephen Rose, Sandra Tuppen, and Loukia Drosopoulou, 'Writing a Big Data History of Music,' *Early Music* XLIII, (2015), 649-60, at 651.

³³ Joan Serrà et al., 'Measuring the Evolution of Contemporary Western Popular Music,' *Scientific Reports* 2, (2012), 521, and Matthias Mauch et al., 'The Evolution of Popular Music: USA 1960–2010,' *Royal Society Open Science* 2, (2015).

³⁴ Such approaches have not been without their musicological critics. See for example R. Fink, 'Big (Bad) Data,' in *Musicology Now* (2013). and Z. Wallmark, "Big Data and Musicology: New Methods, New Questions,"

https://www.academia.edu/6442281/Big_Data_and_Musicology_New_Methods_New_Questions.

³⁵ The Digital Music Lab resulted from a project funded by the Arts and Humanities Research Council (AH/L01016X/1). The project was led by Tillman Weyde, and involved contributions from Samer Abdallah, Mathieu Barthet, Emmanouil Benetos, Stephen Cottrell, Simon Dixon, Jason Dykes, Nicolas Gold, Mahendra Mahey, Mark Plumbley, Dan Tidhar and Daniel Wolff. Further information on the project design is given in T. Weyde, Cottrell, S.J., Dykes, J., Benetos, E., Wolff, D., Tidhar, D., Gold, N., Abdallah, S., Plumbley, M. D., Dixon, S., Barthet, M., Mahey, M., Tovell, A., Alancar-Brayner, A., 'Big Data for Musicology,' in *1st International Digital Libraries for Musicology workshop* (London, 2014). S Abdallah et al., 'The Digital Music Lab: A Big Data Infrastructure for Digital Musicology,' *ACM Journal on Computing and Cultural Heritage* 10, (2017) gives more detailed information on the system and its datasets. The system may be freely accessed via <http://dml.city.ac.uk/vis/>

³⁶ AHRC Research Centre for the History and Analysis of Recorded Music:

<http://www.charm.rhul.ac.uk/index.html>

³⁷ <http://www.mazurka.org.uk/>

³⁸ For a detailed overview of the challenges of inferring pitch levels from different legacy artefacts, see Bruce Haynes, *A History of Performing Pitch: The Story of "A"* (Lanham, Md.; Oxford, 2002), 1-53.

³⁹ *Ibid.*, 360.

⁴⁰ See for example 'Orchestral Tuning', *The Musical Standard*, 26 June 1875, 409-410, or W. H. Stone, 'The Causes of the Rise in Orchestral Pitch,' *Proceedings of the Musical Association* 7, (1880), 99-116.

⁴¹ Haynes, *A History of Performing Pitch*, 363.

⁴² This decade is represented almost exclusively by recordings made by the Virtuoso String Quartet over the 7th and 8th April 1905. It should not be taken as representative of the decade as a whole

⁴³ Haynes notes, for example, that in 1988 the Italian government introduced legislation to reduce overall pitch levels in the country, at the instigation of a number of high-profile opera singers. Haynes, *A History of Performing Pitch*, 363.

⁴⁴ See Carl Stumpf, *Die Anfänge Der Musik, Mit 6 Figuren, 60 Melodiebeispielen Und 11 Abbildungen* (Leipzig, 1911). For an English translation and more recent commentary on Stumpf's work, see *The Origins of Music*, trans. David Trippett (Oxford, 2012).

⁴⁵ For some insight into the historical discomforts ethnomusicologists have felt with the subject of music universals, see the various contributions in *Ethnomusicology*, 15(3), p.379-402 (1971) and *The World of Music* vol. 19(1/2), p.2-141. For a more recent discussion, as well as an acknowledgement of the general paucity of musicological scholarship in this area, see Steven Brown and Joseph Jordania, 'Universals in the World's Musics,' *Psychology of Music* 41, (2011), 229-48.

⁴⁶ Dane L. Harwood, 'Universals in Music: A Perspective from Cognitive Psychology,' *Ethnomusicology* 20, (1976), 521-33, at 531.

⁴⁷ For a comprehensive overview of this work see Justin London, *Hearing in Time: Psychological Aspects of Musical Meter* (Oxford, 2004).

⁴⁸ Richard Parncutt, 'A Perceptual Account of Pulse Salience and Metrical Accent in Musical Rhythms,' *Music Perception* 11, (1994), 409-64, at 409. My thanks to Tillman Weyde for bringing this paper to my attention.

⁴⁹ London, *Hearing in Time*, 31.

⁵⁰ The DML software limits analyses to a maximum of 1000 recordings for those searches that produce larger numbers of results.

⁵¹ The spike observable in the Ghana series for the range 71.96-75.36 appears incongruous, and probably results from the DML software identifying machine noise of some kind (disc revolution, for example) in a batch of recordings.

⁵² Brown and Jordania, 'Universals in the World's Musics,' 243.

⁵³ I am particularly indebted to Daniel Wolff for additional work undertaken with the DML to produce calculations underpinning this case study.

⁵⁴ Nicholas Cook, 'Music as Performance,' in *The Cultural Study of Music*, ed. Martin Clayton, Trevor Herbert, and Richard Middleton (London: Routledge, 2003), 204-14. *Beyond the Score: Music as Performance*.

⁵⁵ Bowen, 'Finding the Music in Musicology: Performance History and Musical Works.'

⁵⁶ *Ibid.*, 450.

⁵⁷ Philip, *Early Recordings and Musical Style: Changing Tastes in Instrumental Performance*, 35.

⁵⁸ These figures again use a beat-tracking approach like that taken in case study 2. Recordings in any given year have been aggregated to produce an average tempo figure. This figure is meaningless in relation to any individual performance, but it is a way of comparing averages over time. The number of recordings available for analysis in any given year varies considerably, with single digit figures in some years—with the obvious methodological risks these imply—up to more than 150 recordings in other years. Grouping these results into five-year bands, to reduce an over-reliance on small numbers of recordings, produced very similar results.

⁵⁹ See, for example, Walter Gray, 'The Classical Nature of Schubert's Lieder,' *The Musical Quarterly* 57, (1971), 62-72, at 62.

⁶⁰ See, for example, Eric Clarke and Nicholas Cook, eds., *Empirical Musicology: Aims, Methods, Prospects* (Oxford, 2004); Tim Crawford and Lorna Gibson, *Modern Methods for Musicology: Prospects, Proposals, and Realities*, Digital Research in the Arts and Humanities (Farnham, England; Burlington, VT, 2009); and Tao Li, Mitsunori Ogihara, and George Tzanetakis, *Music Data Mining* (Boca Raton, 2012).

⁶¹ Rose, Tuppen, and Drosopoulou, 'Writing a Big Data History of Music.'

⁶² *Ibid.*, 651.

⁶³ Franco Moretti, 'Conjectures on World Literature,' *New Left Review* 1, (2000), 54-68. *Distant Reading* (London and New York, 2013).

⁶⁴ *Distant Reading*, 57.

⁶⁵ I am conscious of the paradox between this statement and the third case study presented here, a paradox that arises from the fact that, at present, there are insufficient recordings of music by other composers available for the kinds of analysis undertaken here.

⁶⁶ Andrejevic, Hearn, and Kennedy, 'Cultural Studies of Data Mining,' 384.

⁶⁷ Steven Feld, 'A Sweet Lullaby for World Music,' *Public Culture* 12, (2000), 145–71, at 151-54.

⁶⁸ 'Pygmy Pop. A Genealogy of Schizophonic Mimesis,' *Yearbook for Traditional Music* 28, (1996), 1-35; Hugo Zemp, 'The/an Ethnomusicologist and the Record Business,' *ibid.*, 36-56.

⁶⁹ See Òscar Celma, *Music Recommendation and Discovery: The Long Tail, Long Fail, and Long Play in the Digital Music Space* (Berlin, 2010). Perhaps unsurprisingly, there are those in the industries themselves who argue that in fact the opposite is true: see Alistair Croll, *Music Science* (Sebastopol, CA, 2015), 11.

⁷⁰ See also Rose, Tuppen, and Drosopoulou, 'Writing a Big Data History of Music,' 649.

⁷¹ Moretti, 'Conjectures on World Literature,' 68.