

Richard J. Lewis, Tim Crawford and David Lewis

# Exploring information retrieval, semantic technologies and workflows for music scholarship: the Transforming Musicology project

## Digital transformations

Digital technology has already transformed music. It is well over a quarter of a century since digital recordings, most prominently in the form of the CD, became established as the industry norm. The effects of this transformation are all around us today, and are reflected in a corresponding increase in research and development of digital audio within universities. Yet this has so far had little impact on the discipline of musicology, though musicologists, like everyone else, take for granted the convenience of music downloads and streaming services like Spotify or YouTube, or the easy online access to scanned scores via IMSLP and music-library OPACs. This technology and the developments that made it possible offer new opportunities for exploring new and larger musical repertoires, together with the means to investigate patterns of connection and distribution that potentially could challenge the established canons of the discipline.

In the Transforming Musicology project we are working towards a digital transformation of the discipline of musicology by both applying computational techniques to musicological investigations ourselves and also promoting such techniques more widely amongst the scholarly community.<sup>1</sup> Our approach is ‘user-centric’, in that our main motivation is to advance and support musicology rather than to show off a new application for computing; but to do this we apply state-of-the-art digital technologies. In particular we investigate the possible impact of three

specific kinds of technology: Music Information Retrieval; Semantic Web technologies; and network analysis.<sup>2</sup> This report aims to give an overview of the project’s activities at the end of its second year, and some idea of what we hope to achieve. Since much of the work, especially in the area of interest to readers of *Early Music*, is concerned with establishing a practical infrastructure for the project, we also focus on the concept of ‘workflow’ as a means of capturing the processes of musicological practice, and eventually of its discourse, where semantic technologies (introduced below) play a vital role.

We are doing actual musicological research that exploits technology in three contrasting topics, the first of which is of direct interest to readers of *Early Music*. This considers the relationship between 16th-century instrumental music, particularly lute music, and its vocal models; we provide more details of our infrastructural work in this area below (see ‘Workflows in early music’, below). The second topic (chronologically) concerns Richard Wagner’s well-known compositional use of leitmotifs, how they have been identified and communicated since, and whether and how they are perceived by listeners in performance. We have been fortunate to be able to carry out a unique experiment involving members of the audience for a complete performance of Wagner’s *Ring Cycle* at the Birmingham Hippodrome in November 2014.<sup>3</sup> The very complex and comprehensive set of bio-physical data we gathered from our ten enthusiastic volunteers is

currently undergoing analysis and will be reported in the music-psychology literature in due course.<sup>4</sup> Our third topic explores less formal forms of musical expertise and dissemination as expressed on the web through social networks. The Internet provides publication space for expert performers and listeners—along with casual visitors and malicious vandals—to share their knowledge and commentaries on music of all kinds. Although conventional musicological and ethnographical methods can be deployed to study these communities to some extent, the very scale of the activity and its infrastructure demands the development of new tools and methodologies.

In order to broaden the impact of these technologies to the wider musicological community, and to encourage good practice in their use, Transforming Musicology is itself hosting four mini-projects which were the subject of an open call in December 2013. The selected projects cover application of audio analysis methods to electronic music corpora,<sup>5</sup> digitizing, linking and analysing textual primary sources (mainly concert programmes) which contain information about music,<sup>6</sup> a rigorous approach to gathering new sources of *conductus* texts from the Web,<sup>7</sup> and a study of performance practice in traditional music using audio analysis.<sup>8</sup>

### MIR: Music as ‘information’

The first of the core technologies Transforming Musicology employs is Music Information Retrieval (MIR). Following the practices of text information retrieval, MIR treats digitally encoded music as information-carrying documents.<sup>9</sup> Once extracted, music information can be put to use in various operationalized approaches to music research.<sup>10</sup> Within the MIR field, such research mostly includes music-industry biased topics such as genre and mood categorization, automatic beat, key and tuning detection, automatic transcription, and building music recommendation systems.<sup>11</sup> However, in Transforming Musicology, we explore the operationalization of more musicologically orientated research questions. We are especially interested in understanding how digital data and computational methods can be integrated with and complement existing scholarly methods. In this section we present a little technical

background of MIR and then discuss some of these research questions.

In order to work with a musical document computationally as part of a research process, it must first be encoded in such a way that whatever is regarded as its ‘content’ is preserved. The decision of what constitutes ‘content’ may vary depending on context. Audio recordings, which are effectively pre-encoded as digital data, have so many samples per second (typically 44,100 per second for each stereo channel) that the density of data is considerable, but this dense ‘content’ is not necessarily immediately musically useful. It needs to be reduced into more manageable ‘features’ which are suitable for the task in hand and intended to reflect aspects of the signal that are meaningful in an auditory or a musical sense.

This process is known as audio ‘feature extraction’ and a variety of different audio features have been developed over the years. One promising feature for much musicological work is the ‘chroma’ feature, which summarizes the pitch content of the signal as a probability distribution of each of the standard twelve pitch classes of the Western tonal system, sampled at regular intervals.<sup>12</sup> Once extracted and saved, the features can be used as an index to a database of the recordings that they summarize, allowing the ability to perform search operations, for example, finding occurrences similar to a given audio query, or comparing all the recordings in the database to produce an overall ‘similarity map’ for the whole collection.<sup>13</sup> Within Transforming Musicology, we used this technology in an initial study of Wagner’s leitmotifs, searching recordings of the operas of *Der Ring des Nibelungen* for their occurrences; the results were correlated with those from listening tests carried out with around 70 participants.<sup>14</sup>

Musical notation may also be digitized (as in a Sibelius or Finale file), although neither 16th-century chansons nor Wagner operas are commonly distributed in that form. While countless scanned musical sources are now available online as PDF files, it is very hard to extract musically meaningful content directly from them. To be suitable for MIR operations, musical scores must be encoded in a way that records their musical information directly (often referred to as ‘symbolic’ data). There

is a long history and great diversity of encoding strategies, but the most promising current effort—at least for academic purposes—is the Music Encoding Initiative (MEI), which provides a highly extensible format, allowing different kinds of notation to be represented, and enhanced with editorial and text-critical annotations.<sup>15</sup>

Audio data is relatively easy to capture from commercial or archival recordings (once the legal issues around rights have been dealt with) but encoded symbolic data is typically very difficult to acquire. It requires either time-consuming manual data entry or error-prone automatic transcription via optical recognition. While some commercial packages do a fairly decent job of transcribing very clear, simple scores set in modern notation, working with older printed sources and manuscript sources is much more challenging, and specialized software may be needed for each historical style of music notation. We describe below (see ‘Workflows in early music’) how we use such software in workflows to create large sets of historical musical data directly from images of the printed sources.

Within the digital humanities as a whole, there is much interest in conducting investigations at large scales; what Franco Moretti calls ‘distant reading’.<sup>16</sup> This provides a significant area of impact for musicology. The Centre for the History and Analysis of Recorded Music (CHARM) built a collection of recordings of Chopin mazurkas along with detailed expressive timing information and were able to discover relationships between performers such as teacher/pupil, and to uncover a series of recordings fraudulently attributed to Joyce Hatto but which in fact used re-engineered versions of existing recordings by other pianists.<sup>17</sup> As another example, Michael Scott Cuthbert took a holistic approach to the music of the Italian Trecento—moving the musical ‘margins into the center’—and found that a number of long-held beliefs about that culture, such as the dominance of secular music, are not so well founded, and that much more music came out of centres other than Florence than had been previously thought.<sup>18</sup> Cuthbert’s approach involved analysing a database of the surviving sources, but taking care to include all of those normally considered secondary, or ‘marginal’. The ability to include, in principle, all surviving traces of a given repertory in

an objective analysis, rather than relying on expert intuition for the pre-selection step that human-scale analysis demands, can be seen as one of the main potentials for transformation of the discipline of musicology.<sup>19</sup>

### Semantic web technologies and linked data for musicology

The second core technology in Transforming Musicology is the so-called Semantic Web, which builds on the successful infrastructure of the World Wide Web. The web in its present form links billions of documents published in a very decentralized manner. They are written to be read and understood by people, and generally carry machine-readable information only for the purposes of formatting. To interpret the content of a webpage—as a search engine must, for example—the text of the document must be parsed by a computer as if it were a human reader. The Semantic Web extends the web infrastructure to build a web of documents that are readable by machines, by requiring that at least some of the semantics (or meanings) of those documents be expressly encoded following certain technical standards.<sup>20</sup> In this section we discuss how semantic technologies are being applied in Transforming Musicology and our wider contribution to research in the field.

A pragmatic outcome of research on the Semantic Web has been the set of principles known as Linked Data.<sup>21</sup> Linked Data stipulates that items of interest should be addressable by a Uniform Resource Identifier (URI, essentially a web address), that it should be possible to use the http mechanism (as used by a standard web browser) to retrieve information about those items, that the information thus retrieved should be encoded using established standards, and that the information should embed *links* to other, related items of interest. When these principles are followed, a Web of Data begins to emerge as more and more items are interlinked in what Christian Bizer *et al.* call a ‘single global data space’,<sup>22</sup> just as the links in webpages create the World Wide Web as we know and navigate it.

A key requirement of Linked Data is that the nature of the link between items should be made explicit by declaring its meaning. Part of the standardization work of the Semantic Web includes collecting controlled vocabularies—known as ‘ontologies’<sup>23</sup>—of

such link semantics in order to promote re-use of existing link labels and thus to increase the interoperability of data on the Semantic Web; if we can exchange data using common, agreed labels, it becomes much easier to share our meanings.

Libraries and archives are becoming some of the most important sources of Linked Data for scholars, and a significant effort has been made in developing ontologies for the bibliographic domain.<sup>24</sup> The RISM catalogue has recently started publishing its data following Linked Data principles.<sup>25</sup> For example, the URI <https://opac.rism.info/id/rismid/806047034> represents a particular manuscript source of William Byrd's *Miserere mei, Deus*, and links both to a URI representing Byrd himself and to a URI representing the British Library, where the manuscript is held.

In Transforming Musicology, we are extending the work of an earlier Linked Data project,

SLICKMEM,<sup>26</sup> which links the named composers in the Early Music Online collection (EMO, see below) to authoritative names in the Virtual International Authority File (VIAF) and to composers of recorded music in MusicBrainz,<sup>27</sup> and the places of publication to settlements recorded in DBpedia (the Linked Data publication of Wikipedia).<sup>28</sup> By publishing these connections, we are filling in a small part of the 'single global data space' and improving the discoverability<sup>29</sup> of a small corner of music history. More recently we have been able to extend this resource to embrace detailed programming information from BBC Radio 3's Early Music Show (*illus.1*).<sup>30</sup>

As well as publishing data about musical sources, in Transforming Musicology we are exploring ways of capturing and publishing musicological 'workflows', the processes of musical research. This kind of

**S L o B R** Semantic Viewer: Semantic Linking of BBC Radio

SLoBR Work View

Quando'io penso al martire

a.k.a. Quand'io pens a la martire; Quand'io penso al martir; Quand'io pens' al martire

Contributors: Jacques Arcadelt

Recorded as: <http://www.bbc.co.uk/music/records/nz2j8b>

Explore:

- This work's composer
- All episodes featuring this work
- All episodes featuring this composer
- All episodes featuring contemporary works

Archadelt. 28 CANTO

Quando'io penso al martire Quando'io penso al martire Amor

che tu mi dai grauofo e forte Corro per gir'à morte Corro per gir'à mor-

te Cofi sperando i miei danni finire Ma poi ch'io giung'al passo ij

Ch'è port'in questo mar pien di tormento Tanto piacer ne sento ij

to piacer ne sen to Che l'alma si rinforza ond'io nol pas so Co-

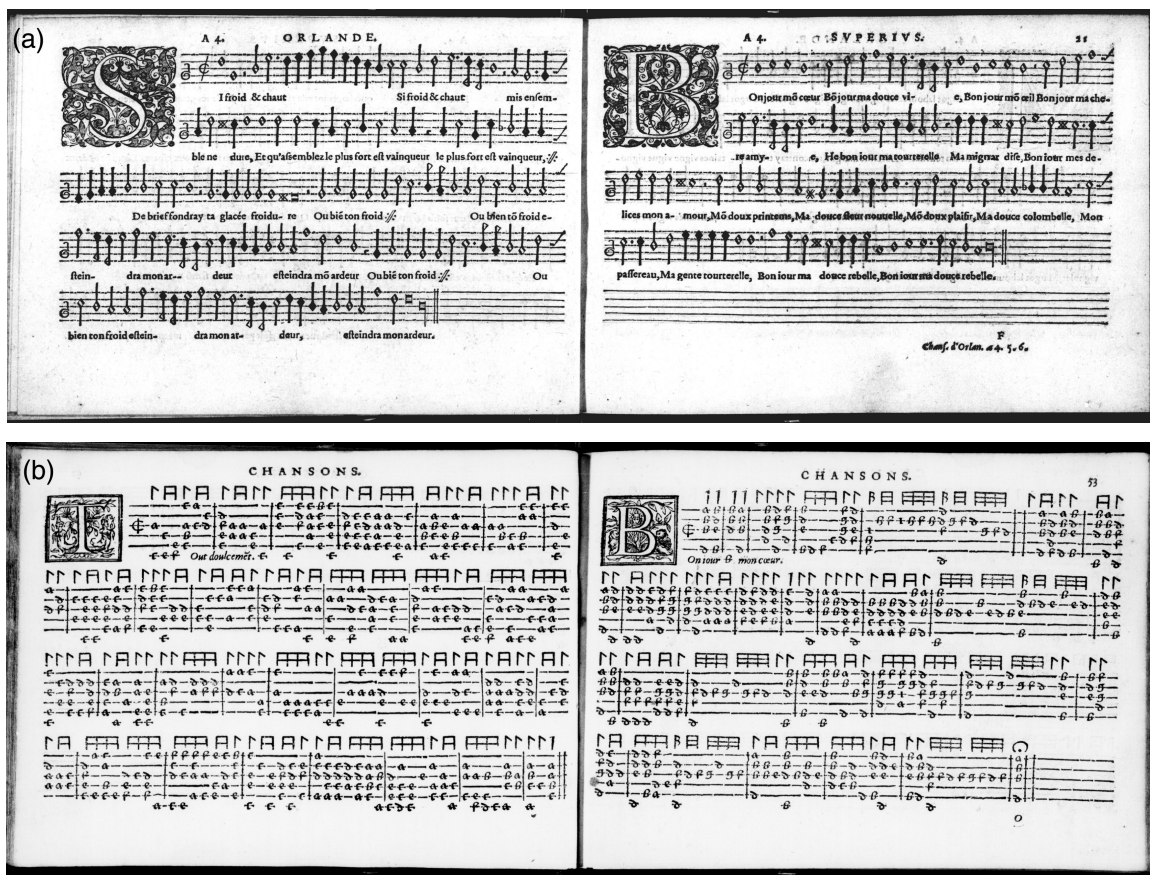
1 A web interface showing how a work found in Early Music Online might be linked automatically to information concerning the music and its composer from Wikipedia or elsewhere, BBC Early Music Show broadcast episodes featuring him, and related episodes of music of a similar period.

work has already been explored in the sciences, where experimental processes and method are already more explicit.<sup>31</sup> While some progress has been made in capturing the workflow of MIR,<sup>32</sup> for musicology our main work has been in capturing workflows around the study of musical materials (i.e. sources), and we describe some of the issues arising around our workflow for automatic lute tablature recognition in the next section. However, much of the process of scholarship—just as in the wider humanities—is really concerned with the construction of discourse through citation and interaction with other primary materials. The challenge for Semantic Web practitioners, and for Transforming Musicology, is

to find meaningful ways of combining automated and scholarly workflows to enhance musical scholarship.

### Workflows in early music

The early music work within Transforming Musicology aims to explore how computational techniques might be brought to bear on questions surrounding the repertory of 16th-century vocal music in its most widely available form as printed partbooks and in its instrumental arrangements, principally for lute. To study this, or any other, area of music history in any detail requires editions of the music and considerable information about its provenance and sources, along with



2 Typical images from Early Music Online (slightly cropped), displaying (a) a single opening from *La fleur des chansons* (Antwerp: Phalèse & Bellère, 1592), RISM B/1 1592<sup>9</sup>, British Library, A.337b, fols.20v–21r; and (b) an opening from *Thesaurus musicus continens ... carmina, ad vsum chelys, vel testudinis accommodata* (Louvain: Phalèse & Bellère, 1574), Brown 1574<sup>12</sup>, British Library, K.1.c.8.(1.), fols.52v–53r. Note the two versions, for voices and solo lute, respectively, of Lassus's chanson *Bon jour mon coeur*, on the recto pages

ways of analysing the connections between pieces, and exploring the musical transformations used by composers and arrangers. Here we build on two earlier projects that worked on methods for producing and curating digital resources of direct interest to scholars and performers of 16th-century music.

Early Music Online is a collection of images of 16th-century printed sources from the British Library, made freely available to scholars and performers as a result of a rapid digitization project.<sup>33</sup> At the same time, the books were newly catalogued, providing much more information about their detailed contents and related scholarly literature; this data is displayed to users of the British Library's online catalogues who encounter EMO materials (see *illus.2a–b*).<sup>34</sup>

The images were digitized from archival microfilms made several decades ago, and represent the British Library's holdings of 300 16th-century anthologies of music (those including music by more than one composer). The total number of musical items is over 8,000, mostly consisting of vocal music

in partbooks. The collection also includes some 30 volumes of tablature, mostly for solo lute.

The Electronic Corpus of Lute Music (ECOLM)<sup>35</sup> aims to make music in tablature more accessible to non-players by providing encodings which can be translated into other notations, into sound, or into formats that can be used for various kinds of analysis and further processing. In this way it intends to support the study of early instrumental music both in its own right and in parallel with the vocal music on which it is most often based. By using web-based interfaces for its software tools where possible, it can support lutenists and others anywhere in the world who are interested to engage directly with their processes. Currently, the ECOLM system contains catalogue information on over 2,000 pieces, of which more than 1,500 are present in full diplomatic encodings which represent the original source as closely as possible.

Most of the encodings in ECOLM were generated by manual input, a slow and laborious process which is particularly prone to error due to operator fatigue. Computers never get tired, so it seemed natural to turn

3 The web-interface used by the ECOLM editorial team for reconciling differences between suggestions made by pairs of correctors, or to override both with a new decision. The ECOLM database stores information about the location of beginnings and endings of pieces, and the process for producing complete pieces is automatic once all systems have been collated in this way

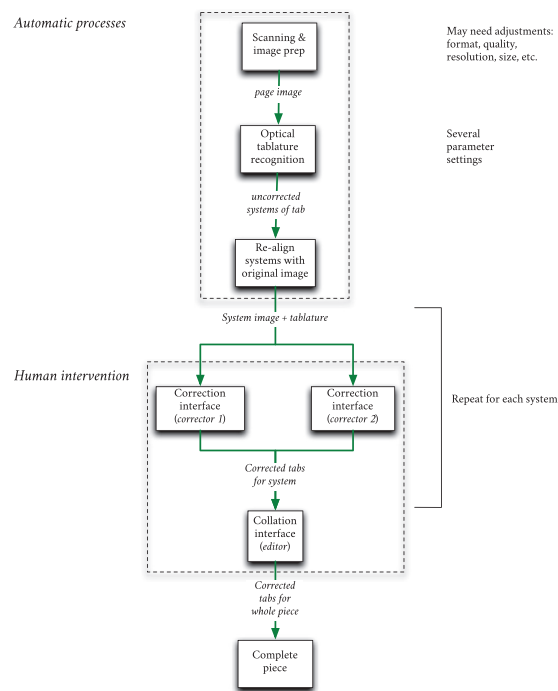
to Optical Character Recognition (OCR) techniques to speed up the capture of data from early printed sources of tablature, since these are usually typeset and use normal type characters, albeit in unusual ways. As commercial OCR programs rely on the use of dictionaries containing words from a specific language (which of course do not appear in tablatures), they are generally unsuitable for the task, but Gamera<sup>36</sup> is a general symbol-recognition package that has the capability of recognizing French, Italian and German lute tablatures.<sup>37</sup> We have been using Gamera's tablature-recognition on the images in books of lute tablature in EMO to generate complete encodings of the books automatically. This is much quicker than manual input, but there remains the problem of inevitable errors, this time introduced by the recognition process, which is highly sensitive to variations in photographic quality or to irregularities in printing caused by damaged type or poor registration. As these can only be put right by human intervention, ECOLM devised a system of double-correction, whereby lines of imperfectly recognized tablature are passed to an online web-based editing program, where they are rendered in alignment with the original image so that corrections can be made quickly and easily in the graphical interface.

For this purpose, lute players are the ones best equipped with the necessary knowledge of 16th-century lute tablature, so ECOLM enlisted the help of members of the UK's Lute Society, about 50 of whom registered to join in this 'crowd-sourcing' exercise. Each line of tablature is corrected twice by different people, whose corrections are identified, logged and time-stamped (see *illus.3*). When every line of tablature in a piece has been double-corrected, another interface is used to highlight differences between the two correctors, allowing a definitive editorial decision to be taken for each case, and this ultimately generates an encoding of the complete piece.

The process just described—from a collection of images to a collection of encoded pieces—is an example of a workflow that can be described formally and recorded in a way that not only ensures that the process is adequately documented but also allows in principle for its repetition (for example with better images, or a newly discovered copy of a source) or its re-use on a different image-set. Throughout the workflow, the settings of the software for each processing stage and the responsibility for each human

intervention are recorded in the management system (*illus.4*). In this way, the complete 'provenance' of the resulting output (the final complete encodings) is saved in such a way that it can be used if necessary at a later stage, either to 'undo' work that has led to a problematic conclusion or, after analysis, to improve the workflow or even the recognition system itself.<sup>38</sup>

The diagram displayed here does not pretend to be fully descriptive of the process (for example, it entirely ignores the back-end functions of the ECOLM database which manages the data flow and maintains the provenance record), but is intended to illustrate some of the elements that a formal description needs to take into account. It is, however, easy to see that the 'human intervention' section of the workflow has a critical effect on the overall process,



4 A (somewhat simplified) view of the workflow for the ECOLM/EMO process of tablature recognition. In the right-hand column are listed some potential 'bottlenecks' in the process, all of which in principle should involve less human supervision as the overall system 'learns' from multiple iterations about the settings that produce best results, the types of correction noticed by the expert correctors and their contexts

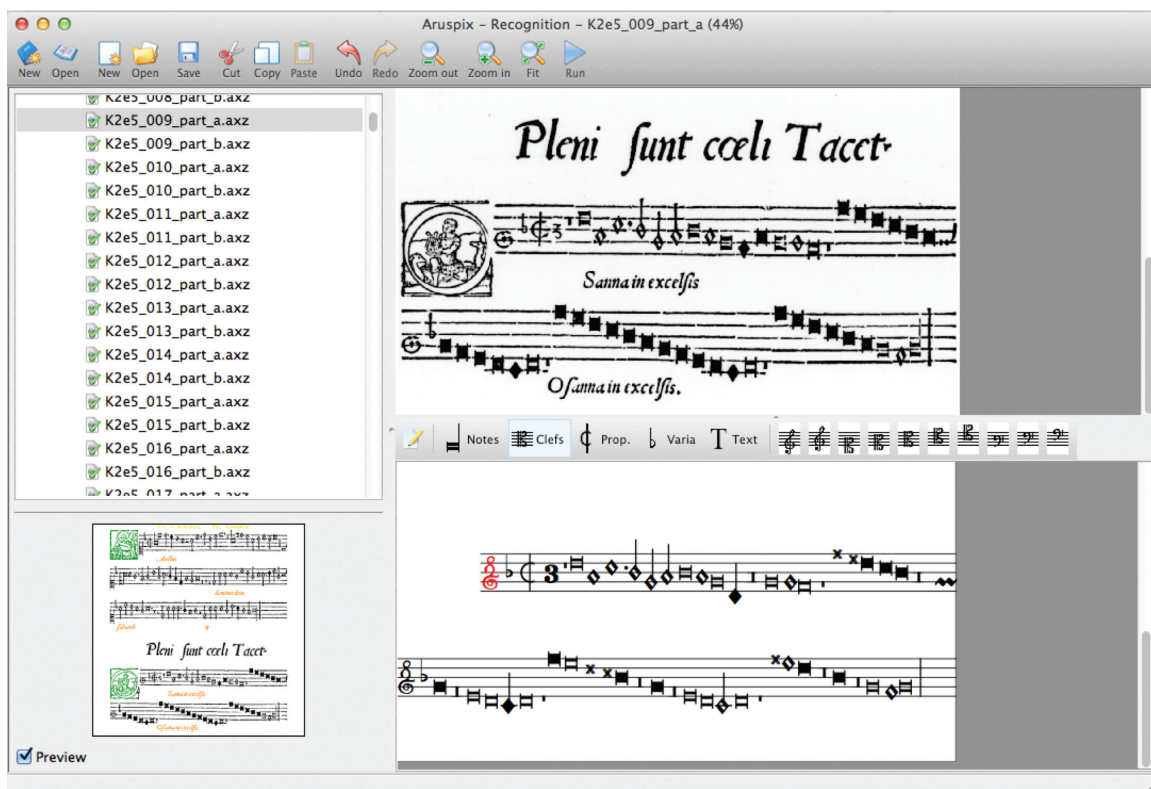
requiring interaction from three people: two correctors, each working on the same individual segments of tablature, and an editor who arbitrates where they disagree. In fact, in a large-scale implementation of this workflow, the corrections themselves can be used to improve the automatic process of recognition in a 'learning' system, at least to the extent that useable encodings can be produced (albeit with a certain level of inevitable recognition error, but much reduced and deemed to be acceptable). After an initial learning phase, human intervention becomes much less necessary.

Another aspect of this workflow not included is the importance of 'metadata', in this case the catalogue data supplied by the library, supplemented by information from the inventories published by Howard Mayer Brown in 1965.<sup>39</sup> This tells the system on which page and on which line of tablature each piece begins. Since pieces sometimes begin midway through a line, the

actual physical location is sometimes required; the way we currently do this is to ask the two online correctors to use the 'New Piece' tool in our correction interface.

A similar process is being applied to the vocal music in EMO, of which the great majority is in partbook form, each vocal part being bound separately but stored by the British Library together with the other parts under a common shelfmark. For the Optical Music Recognition (OMR), we use Aruspix, a software package designed to work with 16th-century printed partbooks.<sup>40</sup>

As can be seen from the screenshot (illus.5), Aruspix has its own graphical editing interface. Corrections made with this are used to improve accuracy by refining the program's internal 'typographical model' (selected to optimize recognition of books printed with a given music typeface) used to perform recognition. As with tablature recognition, the amount of necessary human intervention



5 The Aruspix user-interface for correction of errors in the optical recognition of printed mensural notation, showing part of the Sanctus of Crecquillon's *Missa 'Pis ne me peult'*, from a collection of parody Masses, *Liber primus missarum quinque vocum* (Antwerp: Susato, 1546), RISM B/1 1546<sup>3</sup>, British Library, K.2.e.5.(1.)



declines once detailed corrections have been done on a representative training set of images; thereafter the program is used in fully automatic mode to generate large quantities of encoded music.

Even the best-designed workflow cannot eliminate every potential problem. This report closes with a look at two examples of unanticipated issues that our work has had to devise ways around. The first derives from the fact that the images we used were digitized from archival microfilms, and the second from the relatively underspecified nature of the library metadata.

First, there is the problem of unknown occurrence of duplicate images. Sometimes, openings appear more than once on a film where, for example, the photographer took a second shot to improve focus, or as a result of simple human error. In order to identify

automatically duplicate images in the collection, by combining the results of similarity searches over the music-recognition output from Aruspix with those from a standard image-comparison program, we were able to identify duplicate images very reliably.<sup>41</sup> A problem when using image-comparison programs on these sources is that without building in musical ‘understanding’ they cannot distinguish images of pages from different partbooks which have very similar layout, sometimes dominated (graphically) by the use of large decorative initials (see *illus.6a-b*). Comparing the musical content, however, even when it is imperfect because of OMR recognition errors, improved discrimination greatly.

While the recognition of duplicate images on a film may seem like a trivial task, easily carried out by

(a) *Missa. Morri ma Prius Tho. Crecquillon.*  
 Yrie eleyson ij  
 kyrie eleyson ij  
 kyrie eleyson  
 Hriste eleyson ij  
 Chryste eleyson ij Chryste eleyson  
 Chryste eleyson  
 Yrie eleyson  
 kyrie eleyson ij

(b) *BASSVS. Missa. Pis ne me Peurvenir Tho. Crecquillon. Fo.vi.*  
 Yrie eleyson  
 kyrie eleyson ij  
 kyrie eleyson ij  
 Hriste eleyson Chryste e-  
 leyson Chryste eleyson  
 Yrie eleyson  
 kyrie eleyson ij  
 kyrie eleyson.  
 B: ij.

6 a–b Image-comparison software considered these two pages, from different vocal parts in two different parody Masses by Crecquillon (printed by Susato, 1546; see the caption to *illus.5*), to be rather similar, probably because of the use of relatively large and identical decorated capital letters, and the similar layouts and texts. The music, however, is quite different (although small local similarities can be seen here, too)

a human researcher, it is a good example of a significant issue that arises when working with amounts of data beyond the capacity of individuals to manage. Duplicate images at the very least confuse any process reliant on a predictable sequence of data, such as arose in our second unexpected issue. This concerns the actual 'mise-en-page' of early typeset music, which sometimes shows something of its origins in rivalry with manuscript practices. Contained within the EMO collection are a number of books in choirbook format, where each of (say) four voices are displayed on a page opening. Even where a publication is entirely in partbooks, it is common to find that a single bound partbook contains music for two different voice parts of the same composition on facing pages. Furthermore, just as with lute tablatures, sometimes the voice part for a musical item will begin on the same system as the previous item ends. We need to capture the precise locations of the start and finish points for musical items within each partbook, but the EMO catalogue metadata simply

gives a listing of works and their composer attributions in sequence. Fortunately, at recognition time Aruspix saves the musical content together with the physical location of each symbol on the recognized pages, providing extra information that can be used in post-processing. We use the catalogue sequence as an initial informed 'guess' as to the location of the music for each work in turn, assuming by default that each occupies a single page. Since this is by no means always the case, we provide an easy-to-use online interface for assigning the starts and continuations of musical items on each page, allowing for the fact that these are sometimes not always obvious or logical at first sight, often being pragmatic solutions to the problem of page-turns, for example, just as in manuscripts of the period (see [illus.7](#)).

Once we have the music of a significant and representative collection of 16th-century vocal and lute music in encoded form, we can perform various kinds of search and matching operations across the corpus. It is beyond the scope of this

The screenshot shows the EMO interface for assigning musical items. At the top, there is a toolbar with icons for zooming and editing, and a 'Same as' dropdown menu. Below the toolbar is a grid of musical score images. The grid is divided into four quadrants, each showing a different musical score. The scores are: 12. Quanta belta quanta gratia (T), 13. Poss'io morir di mala morte (T), 14. Ahi se la donna mia (T), 15. Il bianco e dolce cigno (T), 16. Non ch'io non uoglio mai (T), 17. Deh dimm'amor se l'alma (T), and 18. Non ch'io non uoglio mai (T). The list on the right side of the interface contains 24 items, each with a colored circle (S, A, T, B) and a diamond symbol indicating the assigned voice parts.

7 The web-based drag-and-drop interface developed for assigning start- and end-points, and continuations, for musical items in Early Music Online partbook images. By default, the sequence of items from the library catalogue metadata is assigned in turn to each page; in practice, adjustments need to be made, with items frequently 'flowing' from one verso page to the following recto in a manner that minimizes page-turns. In some cases, different voice parts appear on facing pages, and sometimes they continue unexpectedly in a way that renders automatic assignment impossible (such as in the case of a few items that begin on a recto page and continue on the facing verso)

report to discuss the technical details of the algorithms behind these processes, but we intend to use them to reveal new information concerning the use of quotation, reference, paraphrase and allusion of thematic material in the music of the 16th century in its most widely disseminated form. In particular, we wish to focus on the ways

in which music publishers exploited the popularity of certain composers, and even of certain works, and how the amazing range and variety of music they produced for domestic consumption supports or conflicts with the established canon of musicology.

---

*Richard J. Lewis is a research associate at Goldsmiths, University of London. He received both his BA in Music and his MMus in Critical Musicology from the University of East Anglia, and his doctoral work, carried out at Goldsmiths, explores issues around the uptake of computational techniques by musicologists. He currently acts as project manager for the Transforming Musicology project. [richard.lewis@gold.ac.uk](mailto:richard.lewis@gold.ac.uk)*

*In an earlier life, Tim Crawford worked as a professional lutenist, playing on several recordings made during the 1980s. As a musicologist he studies lute music of the 16th to 18th centuries. Since the early 1990s Tim has also been active in the rapidly expanding field of Music Information Retrieval and served as President of its international society, ISMIR, for two years. He is currently Professorial Research Fellow in Computational Musicology at Goldsmiths, University of London, where he is Principal Investigator of the Transforming Musicology project. [T.Crawford@gold.ac.uk](mailto:T.Crawford@gold.ac.uk)*

*David Lewis is Research Fellow at Goldsmiths, University of London and Birmingham Conservatoire. His research focuses on the creation, dissemination and use of digital corpora of music and music-theoretical texts. Projects in which he is involved include Transforming Musicology, the Electronic Corpus of Lute Music, The Complete Theoretical Works of Johannes Tinctoris: A New Digital Edition and Thesaurus Musicarum Italicarum. [d.lewis@gold.ac.uk](mailto:d.lewis@gold.ac.uk)*

*Transforming Musicology is supported by a Large Grant (AH/L006820/1) under the AHRC's Digital Transformations Theme. The research reported here relies heavily on the work of other members of the team, most notably Kevin R. Page and David M. Weigl, both of the University of Oxford eResearch Centre, and Ben Fields (Goldsmiths), whose expertise is crucial to the work in the Semantic Web and Workflows sections above.*

1 [www.transforming-musicology.org](http://www.transforming-musicology.org) (accessed 30 June 2015).

2 The project includes as partners Goldsmiths' College, University of Oxford, Queen Mary University of London, Lancaster University and Utrecht University, and is funded for three years from October 2013 under the AHRC's Digital Transformations Theme.

3 This work (not originally envisaged in the Transforming Musicology project proposal) was made possible by generous extra funding from the AHRC Digital Transformations Theme and from the 2014 Being Human Festival, further supported by the British Academy.

4 For results from a related precursor experiment on the determinants of memory for leitmotifs, see D. Müllensiefen, D. Baker, C. Rhodes, T. Crawford and L. Dreyfus, 'Recognition of leitmotives in Richard Wagner's music: chroma distance and listener expertise', in *Proceedings of the European Conference on Data Analysis* (Bremen, 2014).

5 *Large-scale corpus analysis of historical electronic music using*

*MIR tools: informing an ontology of electronic music and cross-validating content-based methods*, Dr Nick Collins with Prof Peter Manning and Dr Simone Tarsitani (all Durham University).

6 *In concert: towards a collaborative digital archive of musical ephemera*, Prof. Rachel Cowgill (Huddersfield University) with Prof. Alan Dix (Lancaster/Birmingham universities), Dr Rupert Ridgewell (British Library), Prof. Simon McVeigh (Goldsmiths) and Prof. Christina Bashford (University of Illinois, Urbana Champaign).

7 *Medieval music, Big Data and the research blend*, Prof. Mark Everist with Dr Nicholas Gibbins and Dr Gregorio Bevilacqua (all Southampton University).

8 *Characterising stylistic interpretations through automated*

analysis of ornamentation in Irish traditional music recordings, Dr Münevver Kökier (Birmingham City University) with Dr Peter Jančovič (Birmingham), Islah Ali-MacLachlan (BCU), Prof. Cham Athwal (BCU) and Dr Daithí Kearney (Dundalk Institute of Technology).

9 J. Futrelle and J. S. Downie, 'Interdisciplinary research issues in music information retrieval: ISMIR 2000–2002', *Journal of New Music Research*, xxxii/2 (2003), pp.121–31, gives an overview of the field of MIR and its relationship to musicology. M. Leman, 'Systematic musicology at the crossroads of modern music research', in *Systematic and comparative musicology: concepts, methods, findings*, ed. A. Schneider (Frankfurt, 2008), pp.89–115, extends the traditional *trivium* of musicology (historical, music theory and ethnomusicology) to include MIR as the 'engineering approaches' to music research.

10 N. Cook, 'Towards the compleat musicologist?', in *Proceedings of the International Symposium on Music Information Retrieval* (London, 2005), expands the concept 'operationalization' to music research. In the statistical literature, the term is used to mean an instrument for making an un-measurable phenomenon measurable. Cook's usage (and our usage here) is more concerned with understanding the procedures of music research.

11 See C. Liem, 'Music information technology and professional stakeholder audiences: mind the adoption gap', in *Multimodal music processing*, ed. M. Müller, M. Goto and M. Schedl, Dagstuhl Follow-ups (2012), for a discussion of the problematic status of such research within a musicological context.

12 Typically intervals of 200 to 1000 milliseconds are used. G. Tzanetakis, 'Audio feature extraction', in *Music data mining*, ed. T. Li, M. Ogihara and G. Tzanetakis (Boca Raton, FL, 2012), pp.43–74, provides a general overview of audio feature extraction, and

M. Müller, F. Kurth and M. Clausen, 'Audio-matching via chroma-based statistical features', in *Proceedings of the International Symposium on Music Information Retrieval* (London, 2005), give a more detailed description of chroma features. Normally, equal temperament is assumed; for related work recognizing historical temperaments, see D. Tidhar, S. Dixon, E. Benetos and T. Weyde, 'The temperament police', *Early Music*, xlii/4 (2014), pp.579–90.

13 M. O. Jewell, C. Rhodes and M. d'Inverno, 'Querying improvised music: do you sound like yourself?', in *Proceedings of the International Symposium on Music Information Retrieval* (Utrecht, 2010), describes how this process works, using most of the same technology as in Transforming Musicology; C. Rhodes, T. Crawford, M. Casey and M. d'Inverno, 'Investigating music collections at different scales with AudioDB', *Journal of New Music Research*, xxxix/4 (2010), pp.337–48, describes the implementation and provides an evaluation of that software.

14 See Müllensiefen *et al.*, 'Recognition of leitmotives'.

15 P. Roland, 'The Music Encoding Initiative', *Proceedings of the First International Conference on Musical Applications using XML* (Milan, 2002), introduced MEI; see [www.music-encoding.org/](http://www.music-encoding.org/) (accessed 9 June 2015). For recent examples of the use of MEI in early repertoires, see three articles in the previous 'Early music and modern technology' issue of this journal (*Early Music*, xlii/4 (2014)): R. Freedman, 'The Renaissance chanson goes digital: digitalduchemin.org' (pp.567–78); P. Roland, A. Hankinson and L. Pugin, 'Early music and the Music Encoding Initiative' (pp.605–11); J. Stinson and J. Stoessel, 'Encoding medieval music notation for research' (pp.613–17).

16 F. Moretti, *Graphs, maps, trees: abstract models for a literary history* (London, 2005).

17 See N. Cook, 'Performance analysis and Chopin's mazurkas',

*Musicae Scientiae*, xi/2 (2007), pp.183–207. Also see N. Cook, 'Methods for analysing recordings', in *The Cambridge companion to recorded music*, ed. N. Cook, E. Clarke, D. Leech-Wilkinson and J. Rink (Cambridge, 2009), pp.221–45, for a discussion of empirical approaches and objectivity in music research. See N. Cook, *Beyond the score: music as performance* (Oxford, 2014), pp.149–58, for an account of the 'forensic' use of this approach on the 'Hatto' recordings.

18 M. S. Cuthbert, 'Trecento fragments and polyphony beyond the codex' (PhD diss., Harvard University, 2006).

19 An example of such an approach of interest to musicologists dealing with textual and library-catalogue data is The Big Data History of Music, a collaborative project between Royal Holloway College and the British Library recently funded by the AHRC Digital Transformations Big Data Research Programme; see [www.royalholloway.ac.uk/music/research/abigdatahistoryofmusic](http://www.royalholloway.ac.uk/music/research/abigdatahistoryofmusic).

20 See N. Shadbolt, W. Hall and T. Berners-Lee, 'The Semantic Web revisited', *IEEE Intelligent Systems Journal*, xxi/3 (May 2006), pp. 96–101, for an approachable introduction to the foundational technologies and to the philosophy of the Semantic Web. A deeper analysis of the social and philosophical ramifications of the web can be found in L. Floridi, 'Web 2.0 vs. the Semantic Web: a philosophical assessment', *Episteme*, vi/1 (2009), pp.25–37.

21 A summary description of the virtues of the approach can be found at [5stardata.info](http://5stardata.info) (accessed 29 June 2015). For a recent large-scale implementation of Linked Data, involving 3.5 million descriptions of museum objects, see [www.britishmuseum.org/research/collection\\_online/search.aspx](http://www.britishmuseum.org/research/collection_online/search.aspx) (accessed 29 June 2015).

22 C. Bizer, T. Heath and T. Berners-Lee, 'Linked Data—the story so far', *International Journal on Semantic Web and Information Systems*, Special

Issue on Linked Data, v/3 (2009), pp.1–22.

23 G. Antoniou and F. van Harmelen, *A Semantic Web primer* (Cambridge, MA, 2/2008), p.11, explain how the word 'ontology', more normally understood as a branch of philosophy, has been 'hijacked' by computer science and is being used as a singular noun.

24 For example, the Semantic Web in Libraries conference has been running since 2009; see [swib.org/](http://swib.org/) (accessed 10 June 2015).

25 See [opac.rism.info/index.php?id=8&L=1](http://opac.rism.info/index.php?id=8&L=1) (accessed 10 June 2015).

26 Semantic Linking of Information, Content, Knowledge and Metadata for Early Music.

27 [musicbrainz.org](http://musicbrainz.org).

28 The data is published at [slickmem.data.t-mus.org/](http://slickmem.data.t-mus.org/) (accessed 10 June 2015). See also T. Crawford, B. Fields, D. Lewis and K. R. Page, 'Explorations in linked data practice for early music corpora', in *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries* (Piscataway, NJ, 2014), pp.309–12.

29 The term 'discoverability' has become common in library science literature to mean the ease with which information can be found by scholars; see M. A. Somerville, B. J. Schader and J. R. Sack, 'Improving the discoverability of scholarly content in the twenty-first century: collaboration opportunities for librarians, publishers, and vendors', Sage White Paper (2012); online version at [www.sagepub.com/repository/binaries/pdfs/SAGEDiscoverabilityWhitePaper\\_WEB.pdf](http://www.sagepub.com/repository/binaries/pdfs/SAGEDiscoverabilityWhitePaper_WEB.pdf) (accessed 16 June 2015). See N. L. Newcomer, R. Belford, D. Kulczak, K. Szeto, J. Matthews and M. Shaw, 'Music discovery requirements: a guide to optimizing interfaces', *Notes*, lxix/3 (March 2013), pp.494–524, for a discussion within musicology.

30 Semantic Linking of BBC Radio (SLoBR). See D. M. Weigl, D. Lewis, T. Crawford and K. R. Page, 'Expert-guided semantic linking of music-library metadata for study and reuse', in *Proceedings of the 2nd International Workshop on Digital Libraries for Musicology* (Nashville, 2015), pp.9–16. Like SLICKMEM, this work was funded by a sub-award from the EPSRC-funded Semantic Media Network (EP/J010375/1); see [gov.epsrc.ac.uk/NGBOViewGrant.aspx?GrantRef=EP/J010375/1](http://gov.epsrc.ac.uk/NGBOViewGrant.aspx?GrantRef=EP/J010375/1) (accessed 11 June 2015).

31 D. de Roure, C. Goble and R. Stevens, 'The design and realisation of the myExperiment Virtual Research Environment for Social Sharing of Workflows', *Future Generation Computer Systems*, xxv (2009), pp.561–7.

32 See K. R. Page, B. Fields, D. de Roure, T. Crawford and J. S. Downie, 'Capturing the workflows of music information retrieval for repeatability and reuse', in *Journal of Intelligent Information Systems*, xli/3 (2013), pp.435–59, for example.

33 See [webarchive.nationalarchives.gov.uk/20140702233839/http://www.jisc.ac.uk/whatwedo/programmes/digitisation/rapididigi/earlymusicon.aspx](http://webarchive.nationalarchives.gov.uk/20140702233839/http://www.jisc.ac.uk/whatwedo/programmes/digitisation/rapididigi/earlymusicon.aspx) (accessed 11 June 2015).

34 [www.earlymusiconline.org](http://www.earlymusiconline.org) (accessed 29 June 2015).

35 [www.ecolm.org](http://www.ecolm.org) (accessed 11 June 2015). Funded from 1999 to 2006 by grants from the AHRB, and in 2012 by the AHRC.

36 [gamera.informatik.hsnr.de](http://gamera.informatik.hsnr.de) (accessed 29 June 2015).

37 The Tablature Recognition Toolkit (based on Gamera) can be found at: [gamera.informatik.hsnr.de/addons/otr4gamera](http://gamera.informatik.hsnr.de/addons/otr4gamera) (accessed 29 June 2015).

38 RODAN, a generalized framework for Optical music recognition workflow, is described in

A. Hankinson, 'Optical recognition infrastructure for large-scale document analysis' (PhD diss., McGill University, 2014). This forms a central pillar of the Single Interface for Music Score Searching and Analysis (SIMSSA) project based at McGill; see [simssa.ca](http://simssa.ca) (accessed 29 June 2015).

39 H. M. Brown, *Instrumental music printed before 1600: a bibliography* (Cambridge, MA, 1965).

40 L. Pugin, 'Optical music recognition of early typographic prints using hidden Markov Models', in *Proceedings of the International Symposium on Music Information Retrieval* (Victoria, BC, 2006), pp.53–6.

41 C. Rhodes, T. Crawford and M. d'Inverno, 'Duplicate detection in facsimile scans of early printed music', in *Proceedings of the European Conference on Data Analysis* (Bremen, 2014).

## Forthcoming special issues

May 2016 - New light on Thomas Tallis

August 2016 - A celebration of Rameau

February 2017 - Relocating English Music

May 2017 - The Romantic Guitar II

August 2017 - Philippe de Vitry