MUSICOLOGY OF EARLY MUSIC WITH EUROPEANA TOOLS AND SERVICES

Erik Duval[1], Marnix van Berchum[2], Anja Jentzsch[3], Gonzalo Alberto Parra Chico[1], Andreas Drakos[4]

[1] erik.duval@cs.kuleuven.be, Dept. of Computer Science, KU Leuven, B

[2] marnix.van.berchum@dans.knaw.nl KNAW-DANS, Utrecht University, NL

[3] anja.jentzsch@okfn.org Open Knowledge Foundation, D

[4] AgroKnow, GR

ABSTRACT

The Europeana repository hosts large collections of digitized music manuscripts and prints. This paper investigates how tools and services for this repository can enable Early Music musicologists to carry out their research in a more effective or efficient way, or to carry out research that is impossible to do without such tools or services. We report on the methodology, user-centered development of a suite of tools that we have integrated loosely, in order to experiment with this specific target audience and an evaluation of the impact that such tools may have on how these musicologists carry out their research. Positive feedback relates to the automation of data sharing between the loosely coupled tools and support for an integrated workflow. Participants in this study wanted to have the ability to work not only with individual items, but also with collections of such items. The use of search facets to filter, and visualization around time and place were positively evaluated, as was the use of Optical Music Recognition and computer-supported analysis of music scores. The musicologists were not convinced of the value of activity streams. They also wanted a less strictly linear organization of their workflow and the ability to not only consume items from the repository, but to also push their research results back into the Europeana repository.

1. INTRODUCTION AND BACKGROUND

The basic aim of the work presented in this paper is to develop services and tools that leverage content in the Europeana Cloud for researchers in digital humanities [4]. In a first year of experimentation, we focused on content in the Wittgenstein archives at the University of Bergen and the Axiom philosophy group at the VU University Amsterdam [5]. In this paper, we report on experimentation in a second year of the project, where we targeted a research community of musicologists that focus on Early Music.

It is important to note that the Europeana Cloud project has a much wider scope: it is concerned with migrating the backend technology of Europeana to a cloud-based infrastructure. The focus of our work is to demonstrate that this technical development enables new tools and services that make it possible for researchers in digital humanities (in the specific case of the work presented in this paper: researchers in Early Music) to either carry out their existing research in a more effective or efficient way, or to carry out research work that is impossible without such tools and services, at least in practical terms, for instance because it would involve too much manual tedious human labor.

In the early phase of the project, as the cloud-based services are still under development, we investigate this issue of added value by loosely integrating existing tools and services accessing the original Europeana services and other suitable services, and by imitating the workflow of the Europeana research platform, which is still under development.

2. RESEARCH GOAL AND METHODOLOGY

2.1 Research questions

In this paper, we address the following research questions:

1. What are the main problems for digital musicologists whose research focuses on Early Music?
2. How can we address these problems and demonstrate the potential added value of cloud-based tools and services on top of large repositories of content like Europeana for Early Music musicologists?

2.2 Methodology

Our basic methodology is User Centered Design [1]. The users of this iteration were musicologists working on Early Music (up to and including Monteverdi). A small group (5 persons) was selected from within the network of the authors. Besides their focus on Early Music, the musicologists in the group share an affinity with technology, and to a different degree are all involved in applying technology to their research practice.

As designers and developers, we had regular formative evaluation sessions over Skype or Google Hangout with the musicologists. (In fact, this worked surprisingly well and allowed for many more regular meetings than we could have organized in more traditional settings with such a diverse, busy and geographically distributed group of participants.) We also had a face-to-face meeting at the

end of the yearly development cycle, for a more in-depth evaluation (see section 6).

In initial meetings the musicologists discussed with us the workflow, computational tools, and content that they currently use.

It is important to note that the evaluation sessions focused on usefulness and usability-in-the-large, i.e. on whether or not the foreseen tools and associated research methodology would actually be of any substantial added value to the researchers involved. We wanted, more specifically, to find out whether our approach could help them to actually change the way they work, whether such an approach would address problems that they may or may not be aware of in their current way of working, etc. Only to a much lesser extent were we interested in finding out whether the Early Music researchers can carry out their current way of working in a more efficient way with our tools and methodology.

3. RELATED WORK

In the past decades, the musicology community in general has been actively involved in the use and development of digital tools for enhancing musicological research. The scholarly study of Early Music is no exception, focusing on very specific problems from this period of music history, while still making use of generic solutions. The development of encoded music formats has been very important, opening up opportunities for musicologists to make use of and analyze machine-readable scores [18]. Seminal work on music encoding is carried out from the eighties onwards, culminating for now in more recent work on how full digital, critical editions of Early Music could be conceived. Further proof of the affinity of the Early Music community can be found in a special issue of the journal *Early Music* (i.e. Volume xlii (2014), No. 4).

Whereas some research has focused on Optical Music Recognition (OMR) for automated metadata generation [11], we rely on metadata from repositories of musical sources (manuscripts, prints) in Europeana and apply OMR techniques in a later step in order to generate a machine readable music encoding for analysis (see section 5.5). In that sense, the scope and goal of the work presented here is more similar to [6], though we focus specifically on Early Music and a User-Centered Design approach for end user tool design and development (section 2.2). An outcome of this approach is that we provide geospatial and time based visualization of search results, rather than a more conventional list of search results, as used in for instance [11]. In fact, we believe that visual approaches to music access remain underexplored, despite some work like [16] and [19]. Our work is a bit different from this earlier work on visualization in that it focuses more on visualizations based on geospatial and time based characteristics of music rather than on visualizing clusters of related music.

The User-Centered Design approach, which is also central to the work presented in this paper, found its way already in the emerging field of 'digital musicology' [2][3] but our focus is on leveraging the content from large-scale repositories for musicology.

4. MAIN PROBLEMS FOR MUSICOLOGISTS

At the initial stage of our work, we identified the following four core problems for the musicologists in our discussions with them:

1. Difficulty of creating the data and metadata needed: the creation of encoded music scores of Early Music (i.e. 'musical data') is a laborious task, which is often carried out with proprietary software packages not suited for the particular types of music notation from this period. Likewise, the metadata on these scores, their original sources, the composers etc. are locked into paper publications and not easily transformed into digital format.
2. Lack of digital corpora with music scores: there are some repositories with music scores for Early Music, like for example CMME (http://www.cmme.org), ECOLM (http://www.ecolm.org), the Josquin Research Project (http://josquin.stanford.edu) and SIMSSA (http://www.simssa.ca) [6], but they are fragmented and it is tedious and time-consuming to go through the different repositories (each with their own query facilities) and do a systematic search for a particular composer or theme.
3. Information exchange and linking of data when working with different tools: although there are specific tools to process music scores, they do not inter-operate and it is again quite tedious and time-consuming to apply different tools on the same content and then to integrate the results of the different tools.
4. Retrieval and analysis of contextual information about the music scores, from bibliographical and historical databases, like the Oxford Music Online (http://www.oxfordmusiconline.com/) or RILM (http://www.rilm.org).

As will become clear in the remainder of this paper, we eventually succeeded in addressing the 1st, 2nd and 3rd problem listed above.

5. TOOL SETUP

5.1 Introduction

In order to investigate how technology can help the musicologists with these problems, we designed, created, integrated and evaluated a set of prototype tools that extends the toolset we prepared for the philosophers the year before. The complete toolset consists of (see Figure 1):

- Ariadne Finder (section 5.2): this tool, personalized for musicologists, helps researchers search and find content coming from Europeana and other sources in a simple and integrated way - the intent is that this tool addresses problem 2 mentioned above;

- TimeMapper (section 5.3): this integrated tool visualizes the search results from the Ariadne finder on a timeline and an interactive map, in order to enable the musicologists to further filter the content and get a better overview of the different resources found on Europeana (http://timemapper.okfnlabs.org);
- Activity Stream (section 5.4): this service, integrated in all the tools, captures and presents the different actions carried out by the users in their interactions with the tools;
- Aruspix (section 5.5): this is an optical music recognition (OMR) tool which transforms prints of Early Music scores into MEI [13];
- Music21 (section 5.6): this is a Python-based set of tools for analysing music encoded as XML (http://web.mit.edu/music21/) [6].

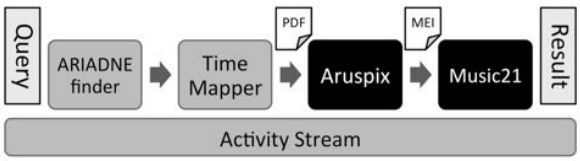

Figure 1: Schema of interconnected tools;

5.2 Ariadne Finder

A series of meetings with the musicology researchers enabled us to identify the content collections of interest. To the Europeana base collection, we added the resources from RISM (http://www.rism.info/) , and integrated them in the Finder. RISM is a well-known and extensively used inventory of musical sources. The abbreviations of library sigla used in RISM, have an authoritative character within musicology, and can be used as a controlled vocabulary in a digital environment.

After the first year experimentation, we simplified the user interface of the Finder by removing some predefined categories from the home screen. Instead, we made a list of four search facets (i.e. provider, media type, language, and year) available on the first screen with the search results.

The integration of the RISM collection was a great challenge: the data covered by RISM (metadata on primary musical sources) are heterogeneous and quite different from the ones provided by Europeana. To allow the integration with the Finder backend and to enable the visualisation of search results in a uniform way, transformation of the metadata to an internal format was required. Moreover, linking to the actual resource was not possible, since RISM provides metadata on the current (physical) holding of the sources, and does not provide links to the digitized versions of the sources.

The Finder is used as the 'baseline' tool for the integration of the other tools, listed below. Both the Activity Stream and the TimeMapper are integrated in the Finder to see the past user activities (i.e. searches) and to visualise search results respectively. When viewing an individual search result, the connection to Music21, through Aruspix, is also available.

In Figure 2, the listing of the search results is shown, with the facets on the left that can be used to further refine the search. Finally, Figure 3 shows how an individual search result is displayed to the user, with the links to the functionality of Aruspix and Music21.

Figure 2 : Search results in the Finder

Figure 3 : Individual search result in the Finder

The Ariadne Finder for the Musicologists group can be accessed at http://greenlearningnetwork.com/cmme-finder/.

5.3 TimeMapper

Europeana provides a variety of metadata for its resources, including thumbnail images, geo-coordinates and time information. TimeMapper visualizes the temporal and geographical characteristics of resources.

TimeMapper is a data visualization tool that allows for the creation of timelines and timemaps using Google spreadsheets (http://timemapper.okfnlabs.org). While the Finder provides the user with a faceted search for Europeana resources, it might still be difficult to navigate through large amounts of search results. We integrated TimeMapper in our tool chain to provide an interactive geo-spatial visualization of the search results. This enables users to quickly navigate the metadata and to order resources on the basis of time and place of publication. In

this way, they can more easily identify resources worth studying in more detail.

Figure 4 shows the TimeMapper when drilling down into resources that match the keyword "Gardano". TimeMapper is available under the MIT licence. The tool can be accessed via the Ariadne Finder button labelled "View in TimeMapper".

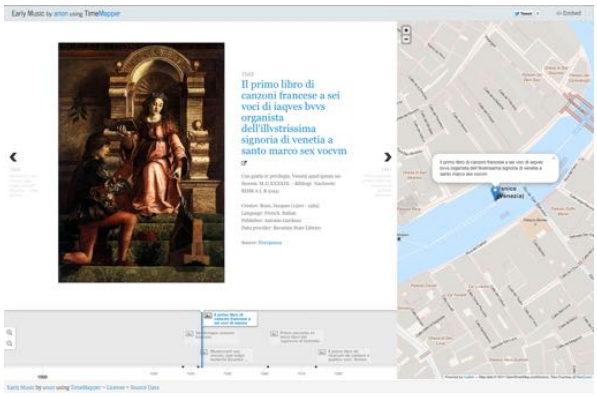

Figure 4: TimeMapper showing resources published by Gardano

5.4 Activity Stream

Based on our earlier work on community reading awareness **Error! Reference source not found.** and supporting the Science 2.0 idea of enhancing collaboration among researchers [17], we have designed, developed and deployed a web application called the "Activity Stream (AS)", enabling researchers to share their work related activities within a community. More specifically in the context of the Early Music musicologists, the application aggregates "search" and "visualize" activities, and makes researchers aware of what their peers are currently working on.

In the first prototype, the AS presented information about "searches" that were carried out with the Ariadne Finder and terms that were "visualized" using the TimeMapper, as illustrated by Figure 5. The activities in the stream are structured as: Actor | verb | (Object). For example, "User from GR" | "has searched" | "Bolzano". For the musicologists, two new activities were added to the activity stream: interpretation and processing. These represent the usage of the Aruspix and Music21 components (see below).

Figure 5: Main screen of the activity stream

The Activity Stream is implemented as a web application (using HTML and JavaScript) and deployed using the Google App Engine (GAE). Together with the terms used to perform a search or visualization, a link to the tool showing the outcome of that action is provided. Also, in order to provide users the flexibility to filter activities, tool grouping was added to the application. For instance, by clicking on the tool's name (e.g.: Finder or TimeMapper) the user can consult the stream of activities from that tool only.

The Activity Stream allows us to digest different events sent from different tools (via REST services) used by researchers, but also provides the possibility to embed these in other software components. For example, the application supports RSS syndication as a passive notification system. Figure 6 illustrates the current activity sources and outlets.

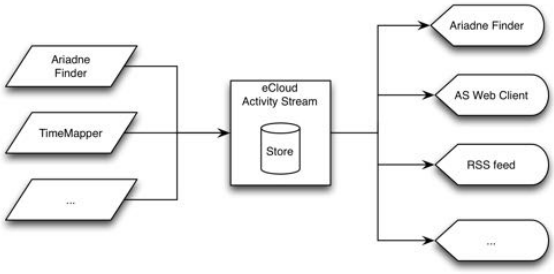

Figure 6: Information sources and destinations of the Activity Stream

5.5 Aruspix

Aruspix is an optical music recognition (OMR) tool that scans early music prints, transcribes them and encodes them into the MEI standard [9][10][15].

While there are other OMR tools available, mainly for music in common music notation, Aruspix is the only tool to our knowledge that can handle scores printed in the 16th and 17th centuries with movable typefaces. Such scores are often difficult to examine with existing superimposition and optical recognition software, as they present a number of specific layout and format problems and are quite often in a deteriorated state because of their age [12][13][14].

The printing techniques of that time mean that differences can exist between copies produced in the same print run, and comparison of these copies by superimposition can enable more accurate critical editions to be prepared. Digitizing the scores through optical recognition can enable us to collate different editions regardless of layout, and is also useful in for instance the preparation of digital music libraries.

For Europeana Cloud, we use the command line version of Aruspix that automatically converts digital scans of scores to MEI files in a page-wise fashion. We then need to combine the pages into a single score again.

Moreover, the MEI version being used by Aruspix is a new and not yet standardized one[13].

Since Music21 (see next section) needs MEI files that use the 2012 or 2013 specification, we developed an XSLT program to transform the MEI files that Aruspix delivers into this newer format.

The command line version sends requested score transcriptions to the Music21 service for further analysis (Section 5.6). Furthermore, it sends activity on transcribed scores to the Activity Stream (Section 5.4).

5.6 Music21

Music21 is a Python-based object-oriented toolkit for computer-aided musicology that allows music information, extraction and generation, together with music notation editing and scripting in symbolic (score-based) forms (http://web.mit.edu/music21/)[6]. The toolkit is able to import different formats, such as MusicXML and MEI.

We extended the Music21 web application module in order to provide parsing and processing requests to a |Music21 installation running on a server. In the workflow, Music21 is used after the Aruspix service has created an MEI version of a score. With an MEI file, a specific set of actions becomes available to the musicologists in order to support them with the analysis of the music involved: calculation of 'Parts and Measures', calculation of the 'Pitch ranges' and requesting the 'legal melodic intervals' of a score.

6. EVALUATION

6.1 General evaluation

To start the discussion, the complete workflow of tools was presented to the musicologists. Afterwards, questions were asked regarding the usefulness of their current tool setup. In general, the participants agreed that the way in which the tools support the research process is helpful. The connection of existing tools (optical music recognition and processing of encoded scores) and automating the process of data sharing between these tools is of great value for them, as it saves them time with their research tasks, compared with using the tools individually. Actually, some of the musicologists had not been able to manually feed the output of one tool as input to the next tool in the workflow.

While the participants found the overall workflow useful, they were also interested in details about specific parts of it. Some of them suggested that, in some cases, just one or two tools are more relevant for their research (e.g. converting a score into a computer readable format or importing their own encoded scores for processing with Music21). This is mainly related to their very varied technical background and research goals. Some of the participants are computational musicologists that regularly use tools like Music21, while others are more traditional musicologists that work with the original sources and have very limited digital research experience.

The participants agreed with the added value of the loosely integrated workflow while doing research on a single item (score), but also observed that the workflow could be automated for use at a larger scale (e.g. a large dataset of scores of a specific period or region). Such automation could be of great value in order to answer research questions about a complete collection or in order to generate new questions for such a collection.

6.2 Ariadne Finder, TimeMapper and Activity Stream

After the musicologists discussed the overall workflow, the loose coupling and setup of tools, they were prompted to assess the tools on an individual level.

From the set of tools adapted from the experimentation the year before with the philosophers [6], the TimeMapper was considered the most interesting and relevant for musicology research. In its current form, the tool provides a visualization of scores based on location and year of print. The participants suggested extending the functionality of the tool, for example with the use of more information than just the data of publication of the prints (e.g. include the information gathered in the Music21 tools, like parallel fifths, valid melodies, or other species counterpoints of a score or measure) or the possibility to compare different timelines that represent results for different search terms. This feedback basically confirms the relevance and usefulness of information visualization techniques in general for musicology research [16][19].

The Finder was mostly seen as a tool that provides existing functionality, similar to what other search engines provide, though the musicologists acknowledged the value of having facets to filter the result set. They suggested to personalize facets to terms that are closer to musicologist research practice, for example, to use 'printed books', 'manuscripts', 'single pieces' instead of 'image' or 'text' classification.

The musicologists were more critical about the usefulness of the Activity Stream (AS) in their research activities. They were not sure that the current actions are relevant for them or even which alternative kinds of activities might be useful to be displayed in the tool. They mostly perceived the AS as an interesting communication device or as a source of information that is comparable to what is common in a Social Network (like Facebook, or more specific for research, like https://www.academia.edu or https://www.researchgate.net/). The participants suggested functionality to enhance the perceived usefulness of the stream, such as a search for specific activities, the possibility to aggregate activities in order to obtain statistics from them, and the possibility to store results for later use.

Participants also suggested other interesting ways to connect the tools, instead of only having a linear approach, as in the current setup. For example, they mentioned that it would be interesting to be able to take the output of Music21 (e.g. parallel fifths of a score) and map

the results, based on their location, with the TimeMapper. This can provide an overview of specific score characteristics and relate them to a particular location.

6.3 Aruspix and Music21

While the Aruspix version included in our workflow does not have a visual frontend for the users, the musicologists acknowledge its importance in the workflow. As mentioned, optical music recognition (OMR) is a crucial step for them [11][12][13][14]. Regarding the current output of this tool, the musicologists would appreciate to see the encoding result and the percentage of errors after the OMR process. While in other sciences, researchers are used to work with and accept a certain percentage of errors, these may not be well accepted in the musicology domain where there is much less of a tradition to work with data that include errors. Nevertheless, the musicologists appreciate what is happening behind the scenes and how good the obtained encoding is, and believe that the results could build trust from the user in the system. Moreover, information about errors can be used as a feedback mechanism for Aruspix: study participants mentioned that they wanted such a facility to be as simple as possible but at the same time complete enough to get the desired information.

The Music21 web interface was one of the most interesting tools for the musicologists. Besides the textual rendition of the analytical results, the participants would also like access to plots or statistics (e.g. note distribution), as these could be more helpful in order to identify characteristics of a score. Currently, the Music21 interface only supports a specific set of generic calculations and processes [6]. The participants would like to have the freedom to build their own analysis, via text or through a graphical user interface.

6.4 Other comments

During the face-to-face evaluation session, the participants provided suggestions about the tools and the workflow, but also about the underlying concepts. For example, some users suggested being able to push the generated encoded scores by Aruspix (MEI or MusicXML) back into the Europeana repository, so that we would use OMR technology to generate metadata, as in [12][14]. Likewise, results created with the Music21 toolkit could be considered as metadata of a particular composition, and as such could also be fed back into the Europeana Cloud repository. Such an approach would enable sharing intermediate research results with peers and a more Science2.0 approach to research [17].

While it was not the direct scope of our work, the participants made a number of suggestions for enhancing the specific usability of the tools and providing a nicer user interface overall.

Finally, the participants suggested additional tools or functionality to be considered. These included:

- Possibility to run batch processes, in order to get a broader overview of music characteristics of a set of scores.
- Support for playback mechanisms in Music21 (or Aruspix), in order to be able to validate and confirm the automatic encoding by listening to the result.
- Possibility to annotate directly into the digital version of a score.
- Possibility to create their own visualizations based on the data obtained from different tools, especially from the Music21 output.
- Inclusion of additional musicology resources, for example from http://www.diamm.ac.uk/.

7. CONCLUSION AND FUTURE WORK

Basically, the User-Centered Development process seems to work as intended: the target users positively evaluated the end result. An important issue for the next cycle is to connect the frontend tools for researchers with the actual backend infrastructure of Europeana Cloud, which has progressed into deployment while our work was taking place. This integration in the production system will enable us to work with more comprehensive content collections.

It is clear from the results that we obtained that there is substantial potential to support novel research methods on large-scale collections of music sources, using technologies like Optical Music Recognition, information visualization, loose coupling of tools, and flexible search. Our work illustrates how this can help researchers in Early Music to carry out existing research in more efficient and effective ways, and even address research questions that are hard or impossible to work on with more traditional means. As such, the potential for a Science2.0 approach to musicology is quite considerable.

8. ACKNOWLEDGEMENT

We gratefully acknowledge the support of the Europeana Cloud project, funded under the ICT Policy Support Programme (ICT PSP) as part of the Competitiveness and Innovation Framework Programme by the European Community (grant agreement no. 325091).

We are very grateful for the intensive discussions and detailed feedback from the following Early Music musicologists: Frans Wiering (Utrecht University), Reinier de Valk (City University London), Eliane Fankhauser (Utrecht University), Laurent Pugin (RISM Switzerland) and Peter van Kranenburg (Meertens Institute - KNAW)

9. REFERENCES

[1] C. Abras, D. Maloney-Krichmar, and J. Preece, ”User-centered design,” *Encyclopedia of Human-Computer Interaction*, Vol. 37, No. 4, pp. 445–56, 2004.

[2] A. Aljanaki, D. Bountouridis, J. A. Burgoyne, J. Van Balen, F. Wiering, H. Honing, and R. Veltkamp, “Designing games with a purpose for data collection in music research. Emotify and hooked: Two case studies,” *Games and Learning Alliance*, pp. 29–40, 2014.

[3] M. Barthet and S. Dixon. “Ethnographic observations of musicologists at the British Library: Implications for music information retrieval,” *ISMIR11: Proc. of the 12th International Society for Music Information Retrieval Conference*, pp. 353–358, 2011.

[4] A. Benardou, C. Dallas, and A. Dunning, “From Europeana Cloud to Europeana Research: The challenges of a community-driven platform exploiting Europeana content,” *Digital Heritage. Progress in Cultural Heritage: Documentation, Preservation, and Protection*, pp. 802–810, 2014.

[5] G. Parra, J. Klerkx and E. Duval, "What Should I Read Next?: Awareness of Relevant Publications through a Community of Practice," C*HI '13 Extended Abstracts on Human Factors in Computing Systems*, pp. 2375-2376, 2013.

[6] H. van den Berg, G. Parra, A. Jentzsch, A. Drakos and E. Duval, "Studying the History of Philosophical Ideas: Supporting Research Discovery, Navigation, and Awareness," *Proc. of the 14th Inter. Conf. on Knowledge Technologies and Data-driven Business*, 12:1-12:8, 2014.

[7] M. S. Cuthbert and C. Ariza. “Music21: A toolkit for computer-aided musicology and symbolic music data,” *ISMIR10: Proc. Of the 11th International Conference on Music Information Retrieval*, pp. 637–642, 2010.

[8] I. Fujinaga, A. Hankinson, and J. E. Cumming, “Introduction to SIMSSA (Single Interface for Music Score Searching and Analysis)”, *DLfM14: Proc. of the 1st International Workshop on Digital Libraries for Musicology*, pp. 1–3, 2014.

[9] A. T. Geertinger, and L. Pugin: “MEI for bridging the gap between music cataloguing and digital critical editions”, *Die Tonkunst. Magazin für klassische Musik und Musikwissenschaft*, 5.3, pp. 289-294, 2011.

[10] A. Hankinson, P. Roland, and I. Fujinaga, “The Music Encoding Initiative as a document-encoding framework”. *ISMIR11: Proc. of the 12th International Society for Music Information Retrieval Conference*, pp. 293–298, 2011.

[11] A. Hankinson, J. A. Burgoyne, G. Vigliensoni, A. Porter, J. Thompson, W. Liu, R. Chiu, and I. Fujinaga, “Digital document image retrieval using optical music recognition,” *ISMIR12: Proc. of the 13th International Society for Music Information Retrieval Conference*, pp. 577–582, 2012.

[12] L. Pugin, ‘Optical Music Recognition of Early Typographic Prints using Hidden Markov Models’, *ISMIR06: Proc. Of the 7th International Conference on Music Information Retrieval*, pp. 53-56, 2006.

[13] L. Pugin, J. Hockman, J.A. Burgoyne, and I. Fujinaga, “Gamera Versus Aruspix: Two Optical Music Recognition Approaches,” *ISMIR08: Proc. of the 9th International Society for Music Information Retrieval Conference*, pp. 419-424, 2008.

[14] L. Pugin and T. Crawford, “Evaluating OMR on the Early Music Online Collection,” *ISMIR13: Proc. of the 14th International Society for Music Information Retrieval Conference*, pp. 439–44. 2013.

[15] P. Roland, A. Hankinson, and L. Pugin, "Early music and the Music Encoding Initiative," *Early Music*, Vol. xlii, No. 4, pp. 605-611, 2014.

[16] M. Schedl, P. Knees, K. Seyerlehner, and T. Pohle, “The CoMIRVA Toolkit for Visualizing Music-Related Data”. *EUROVIS07: Proc of the 9th Eurographics/IEEE-VGTC Symp. on Visualization*, pp. 147-154, 2007.

[17] B. Shneiderman, “Science 2.0,” *Science*, Vol. 319, No. 5868, pp. 1349-1350, 2008.

[18] J. Stinson and J. Stoessel. “Encoding medieval music notation for research,” *Early Music*, Vol. xlii No. 4, pp. 613–617, 2014.

[19] S. Stober and A. Nürnberger. “A multi-focus zoomable interface for multi-facet exploration of music collections,” *7th International Symposium on Computer Music Modeling and Retrieval*, pp. 339–354, 2010.