

Laurent Pugin

Aruspix and the Marenzio Online Digital Edition: Some Lessons and the Evolution of the Project Plan

1. Introduction

Critical edition projects are complex and notoriously long endeavours, and music is no exception. The publication of a critical edition of the complete work of a composer typically takes decades and represents a significant amount of scholarly work. Very often, the project needs to be adjusted in the light of new findings in the field. New sources may be discovered, authorship of pieces or manuscripts may be confirmed or oppositely invalidated. In some cases, this can yield the necessity to revise or extend significantly the editorial criteria for the project. Over time, editorial and music performance practices can change, which is often but not exclusively the case with Early music. For example, reducing note durations for mensural notation, which used to be the standard practice a few decades ago, is nowadays not the rule anymore. As a result, editorial projects often have to balance the tension between a desired uniformity in its outcome and the expectation to comply with state-of-the-art and most recent developments in the field of musicology and editorial practices.

In most countries, there are no dedicated funding schemes for critical music editions which makes it very difficult to finance them. The notable exception is Germany, where many leading music edition projects benefit from long-term funding, mostly through the *Akademie der Wissenschaften und der Literatur* in Mainz. The *Akademie* funds several complete edition projects.¹ For countries that do not benefit from this unique type of funding, critical edition projects are usually supported directly through the involvement of university collaborators or through dedicated research projects. However, raising support for editorial projects through research funding schemes is usually quite difficult. Preparing critical editions is not *per se* research, or at least is often not considered as such. Furthermore, the typical duration of a research project (three or five years) remains a relatively short timeframe for the realisation of a proper editorial venture compared to the usual needs of a complete work edition project.

¹ There are currently 16 long-term musicology projects related to music editions and music documentation supported by the *Akademie* for a total of 86 collaborators.

The development of digital technology adds another layer of complexity that comes with an even more rapid evolution pace. Technology evolves extremely quickly, and new possibilities are constantly emerging. In the last decade, the field of digital critical edition has established itself, first as a research field and now as a new direction for scholarly editing and publishing.² Digital editions offer alternative ways to manage the evolution of the editorial project and present several advantages over traditional printed editions. One of the most important is the possibility to revise those parts of the editions that have already been published and to react to the possible findings and needs for changes mentioned above. This can, to some extent, include the revision of the editorial criteria, even though doing so could yield significant revision work that obviously ought to be avoided as much as possible.

As always, innovation and technological developments come with new challenges and new questions. Although revising and adjusting already published content throughout the editorial project is to be considered a great advantage of digital editions, the direct drawback is that this requires the appropriate infrastructure and workflows for making these adjustments. This is not necessarily straightforward, simply because editors who have finished and published a volume are not always available or ready to make the desired revisions. Updating content requires appropriate schemes to be put in place for users to be able to properly cite and access the edition for each specific publication stage.

These broad considerations concern only some of the transformations that digital technology can bring to the progress of an editorial project but does not touch upon all the possibilities and challenges brought by digital editions regarding user interfaces, navigations or any other aspects. Furthermore, the development of digital technologies can also have some more in-depth consequences in the evolution of an editorial project. These are the topic of this paper which provides an example of the significant changes the Marenzio Online Digital Edition (MODE) project went through over time, adjusting itself to the very fast developments the field of digital music editions went through over the past fifteen years.

2 Johannes Kepper and Laurent Pugin, »Was ist eine Digitale Edition? Versuch einer Positionsbestimmung zum Stand der Musikphilologie im Jahr 2017,« *Musiktheorie, Digitalität in der Musikwissenschaft* vol. 32 no. 4 (2017), pp. 347–63.

2. History of the project

The Marenzio edition project started in 2005 under the direction of Mauro Calcagno, when music digital editions were still in their infancy.³ This editorial project was launched with two international conferences held in 2005 and 2006 at the Accademia Nazionale di Santa Cecilia in Rome and at Harvard University respectively. The main goal of the two conferences was to assess the current state of research on the composer and to lay the groundwork for the new edition. The conferences generated two separate publications, which include five articles devoted specifically to the Marenzio edition.⁴ During the conference in Rome, preliminary discussions took place for integrating digital technology within the editorial process with the use of Aruspix, a tool developed at the University of Geneva under the supervision of Etienne Darbellay. The main idea of Aruspix was to offer a suite of tools for comparing printed music sources of the sixteenth and seventeenth century, similar to what had been developed for text with, for example, the collecting and exhibiting software Collex and the textual collation tool Juxta used in the NINES (Nineteenth-Century Scholarship Online) project.⁵

With this approach, the project very quickly took a direction that went beyond the benefit of making a corpus of music more readily available. Indeed, the idea of the project became to offer a unique link between the world of the Renaissance printing press and that of today's digital technologies. The technology of printing by movable type, which in the literary domain transformed European culture and was used to publish all of Marenzio's works, was the most common way of printing music in the sixteenth century. Despite the fact that it was cheaper and that it generated a profitable market, particularly in Venice (where most of

3 At that time, digital technology related to music notation publishing was mostly embedded in music notation software applications that were designed to produce printed editions, such as the Wolfgang music notation software application. The only notable exception was the Computerized Mensural Music Editing (CMME) project focusing on Renaissance music, with a Java-based application plugin that would run in web-browsers once installed by the user. See CMME, <https://cmme.org>. (All URLs in this paper were accessed in October 2019.)

4 They are, in *Luca Marenzio e il madrigale romano*, ed. Franco Piperno (Rome: Accademia Nazionale di Santa Cecilia, 2008): Mila De Santis, »Problemi di edizioni di testi poetici intonati nel Cinquecento (con alcuni esempi marenziani),« and Maria Caraci Vela, »Per una moderna edizione critica della musica marenziana« ; and in *Perspectives on Luca Marenzio's Secular and Devotional Music*, eds. Mauro Calcagno and Paolo Cecchi (Turnhout: Brepols, 2011): Etienne Darbellay, »Do We Need a New Edition of Marenzio's Music?,« Christine Jeanneret, »Dare in luce et con diligenza correggere: A Study of Marenzio's Editions and Re-editions,« and Laurent Pugin, »Music Printers at Work: Comparing Editions of Marenzio's *Primo libro di madrigali a quattro*.« The latter volume was published as part of a series under the auspices of the Centre d'Études Supérieures de la Renaissance in Tours, directed by Philippe Vendrix.

5 NINES, <http://www.nines.org>.

Marenzio's works were published), this technology presented several limitations that became especially prominent during Marenzio's time: the text setting, for example, was often misaligned with the music, precluding a clear reading. Because of Marenzio's popularity and the numerous editions and reprints of his works, the objective of the project was to present the best-case scenario in which to study the dynamics of the music printing market in the late Renaissance. As the digital tool suite of Aruspix was developed in conjunction with the Marenzio edition, it would facilitate an investigation of these dynamics with unprecedented depth, while producing a modern edition usable by scholars as well as professional and amateur musicians.

The Marenzio edition project has not been supported by a long-term funding scheme. However, it has been continuously supported by the institutions of the scholars involved in it. In the first place, the project received support from the various university libraries.⁶ The project received dedicated funding for the editorial work through the Swiss national science foundation, a fellowship award to Christine Jeanneret from 2009 to 2011,⁷ and most importantly through a three-year grant (2011-2014) from the National Endowment for the Humanities (NEH) in the U.S.A. as part of the »Scholarly Editions and Translations« program awarded to Mauro Calcagno (PI) and Giuseppe Gerbino (Co-PI).⁸ The project also received regular funding from institutions for workshop and additional editorial work.⁹

3. Evolution of the project

At the time of the NEH project in 2011, the Marenzio edition had already evolved from the original idea to producing a traditional paper-based edition to creating a digital edition. The project became the MODE project even though not all technology challenges had been solved at this time. During these years, the project

6 It includes the Isham Memorial Library of the Eda Kuhn Loeb Music Library at Harvard University, the Stony Brook University Libraries, the Gabe M. Wiener Music & Arts Library at Columbia University, and the Kislak Center as well as the Schoenberg Institute for Manuscript Studies at the University of Pennsylvania.

7 Marenzio Project: for a Critical Edition of Luca Marenzio's Profane Music, <http://p3.snf.ch/project-121444>.

8 Online Edition of the Secular Music of Luca Marenzio, <https://securegrants.neh.gov/publicquery/main.aspx?f=1&gn=RQ-50553-11>.

9 It includes the Center for Digital Research and Scholarship at Columbia University (CDRS), which has provided technical assistance both for digital storage and for the building and maintenance of the website, the Faculty in the Arts, Humanities, and lettered Social Sciences (FAHSS) initiative at Stony Brook University, and the Price Lab for Digital Humanities at the University of Pennsylvania.

had to go through some quite significant adjustments, influenced by other projects in the field and by the development of new technologies. We can list at least three significant developments that, at the beginning of the project, did not exist or that were not yet very developed and which would have a direct impact on the way the MODE project evolved and adjusted its objectives.

3.1 The development of the Music Encoding Initiative (MEI)

The fundamental question for a digital edition is how to encode the data. The encoding scheme has some direct implications for current and also future interoperability of the project. It determines the ways in which the data will be usable by tools and how it can be shared with other projects. It is also directly related to its long-term sustainability and other long-term archiving questions. For music edition projects, having a common scheme makes it possible to share data between them but also to share tools.

At the beginning of the project, MEI was still at a very early developmental stage. In 2007, however, thanks to the involvement of scholars from the University of Detmold / Paderborn, the project grew significantly and rapidly. MEI clearly was the best candidate for MODE.¹⁰ First of all, MEI includes a very rich metadata header that makes it possible to include very detailed metadata in a highly structured manner. MEI also has all the benefits of any XML format by default. Standard tools for manipulating XML can be used, and it is both machine-readable and at the same time, to some extent, human-readable. One strength of MEI that made it particularly interesting for both the Aruspix project and the MODE project is its flexibility for modelling various types of music notation, including mensural notation. It offers a dedicated module specifically designed for accommodating the special needs of mensural notation. This means, for example, that a minim will be encoded as such and not using its Common Western Music Notation (CWMN) equivalent, a whole note. Such distinctions are important and are one of MEI's advantages.¹¹

MEI also provides an encoding organisation by parts (as opposed to a score organisation), which makes it perfectly suitable for encoding part-books. In a

¹⁰ See MEI, <http://music-encoding.org>.

¹¹ On that base, it is possible to develop tools that are aware of the underlying notational system and that can process mensural note durations appropriately, i.e. to have algorithms that understand the context given the 16 mensural species provided by the four levels of division in mensural notation: *modus major*, *modus minor*, *tempus* and *prolatio*. It means algorithms that are ternary-based (and not binary-based as in CWMN) can be developed. They themselves apply imperfection or alteration rules when the level is imperfected or altered without incorporating in the encoding of the note itself the number of sub-units it represents. In other words, calculating the actual duration of a particular note can be left to the algorithm and does not have to be hard-coded.

part organisation, each part is encoded in a distinct XML sub-tree. With this, it is possible to have all the parts in a single file even if the score has not yet been assembled. For an editorial project such as MODE, however, the question of how to move from a part encoding resulting from an Optical Music Recognition (OMR) process to a complete score encoding raises some important questions. What link to the original source do we want to keep (or can we keep) once we have switched to a score representation? Can we create a diplomatic score transcription of the part-books?

Even if we aim at building a diplomatic score transcription, encoding the score creates a distance from the sources in part-books, since the organising format is no longer the same. As an example, keeping ligatures in a score is cumbersome, and scholars traditionally use artifices such as brackets to mark them in the score.

On a broader scope, these questions illustrate the multiple encoding viewpoints with which this type of music source can be approached, and, above all, how the viewpoint can be expected to change throughout the editorial process. Figure 30 illustrates the various steps in such a process, from the raw OMR output to a CWMN score usable by common sheet music viewers or music analysis frameworks. Each step in this process requires a change of perspective on the music source.

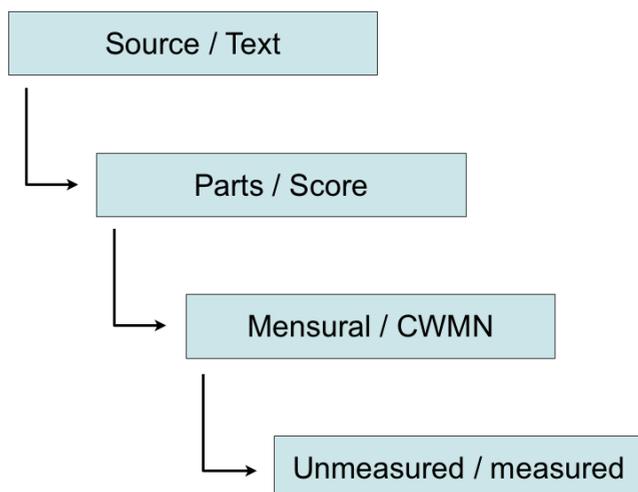


Figure 1. The various steps in a digital edition workflow for early music part-books. Each step requires the encoding viewpoint to be shifted.

Each step of the process presents a duality of possible viewpoints on the data. First of all, the raw output of the data as generated by an OMR process needs to be transformed from a purely source-focused representation (e.g., a note is represented by its staff-line position) to a text-focused one (a note is represented by its pitch). The next step in the process requires the shift from a part-based representation to a score-based one, with all the voices aligned. This requires the aforementioned question of the durations to be resolved. At this stage, a possible viewpoint change lies in moving from a mensural representation with note durations represented by a ternary-based system to a CWMN representation with durations represented in a binary-based system. Finally, the encoding can be transformed from an unmeasured to a measured representation.

A possible approach to tackling this problem has been proposed in experiments for introducing a layout module in MEI. The general idea of the layout module is to separate more effectively the presentation and the content.¹² With the layout module, there is a single XML sub-tree for the content (the usual MEI sub-tree with the actual music encoding) and in addition distinct sub-trees for each of its presentations. The elements in the presentation (or layout) sub-tree always point to the content. Concerning the part-books, each individual part is only one presentation of the content, while the full score is another. With this experimental layout module, MEI is being used in Aruspix with promising results. However, while this approach clearly shows advantages by avoiding data redundancy, future work on this will be necessary to determine if this is a valid option for the MODE project and for similar editorial projects.

Working on content and presentation separation is one example of how collective research efforts have contributed to modifying and improving the MEI schema. It often requires extensive community discussions, which are at the core of the MEI development and make it strong and unique. However, this comes with the cost of the time for these discussions to be conducted and the consensus to be found. Nonetheless, the development of MEI made it possible for the MODE project to become part of a broader community and to establish itself as one of its players.

¹² Laurent Pugin, Johannes Kepper, Perry Roland, Maya Hartwig and Andrew Hankinson, »Separating presentation and content in MEI,« in *Proceedings of the 13th International Society on Music Information Retrieval Conference (ISMIR 2012)*, Porto, Portugal (2012), pp. 505–510.

3.2 The development of Verovio

In 2013, the Swiss RISM Office launched the development of an open-source software library for rendering music incipits named Verovio.¹³ The main idea behind the development of Verovio was to implement a tool that could render MEI natively, i.e. without converting MEI either explicitly or internally in the software application used for rendering to another format. With Verovio, the MEI code is parsed and rendered as notation with a single tool and in one step. Verovio has been developed as a software library and not as a full software application. This means that it is not a desktop music notation application but instead a software component that can be integrated into a wide range of application environments.

The decision was made to develop Verovio from scratch in order to be able to operate on a memory representation of MEI. Verovio directly implements the MEI structure internally with the exception of a top-level page-based organisation that is required for the organisation of the rendering. The reason for choosing to implement a library from scratch rather than modifying an existing library such as GuidoLib¹⁴ is that in the long run it will make it significantly easier to render complex MEI features. Previous experience has indeed shown that modifying an existing solution can be very quick to develop at the beginning but that the development curve eventually reaches a plateau.

Another idea behind the development of Verovio was to have a tool that would produce SVG in order for it to be easy to use in web environments. With the same goal in mind, it was designed to be light, fast, and with no external dependencies, making it very flexible and easy to use or integrate into a wide range of environments. This opens up a whole range of different possible uses. A JavaScript version of Verovio is particularly promising because it provides a fast in-browser music MEI typesetting engine that can easily be integrated into web-based applications. This setup makes it possible to design ground-breaking web applications where the MEI encoding is rendered on the fly. In such designs we can rethink the interface and avoid mimicking page output. We can instead adjust the layout dynamically to the screen of the device employed by the user.

13 Verovio is named after the Dutch music engraver, Simone Verovio, who was active in Rome in the late sixteenth century. He was one of the very first to use engraving on copper plates for printing music, a technique that was revealed to be the most suited to music about one century later. See Laurent Pugin, Rodolfo Zitellini and Perry Roland, »Verovio: A library for engraving MEI music notation into SVG,« in *Proceedings of the 15th International Society for Music Information Retrieval Conference (ISMIR 2014)*, Taipei, Taiwan (2014), pp. 107–12. See also Verovio, <https://www.verovio.org>.

14 Christophe Daudin, Dominique Fober, Stéphane Letz and Yann Orlarey, »The Guido Engine: A toolbox for music scores rendering,« in *Proceedings of Linux Audio Conference*, Parma, Italy (2009), pp. 105–111.

The layout can be calculated to fill the size of the screen or interactively changed according to a zoom level adjusted by the user.

When the project started, some music digital edition projects were using CD-ROM or USB-keys as publication media, storing desktop software applications that users had to install on their computers. It quickly became clear that these solutions would be difficult to maintain and that a web-based solution would be more desirable. For MODE, the development of Verovio filled a crucial gap and made a web-based edition possible. It was the missing piece for putting in place a truly dynamic online music edition that would not require to distribute a software application which users would need to install locally. This was even more of a break-through considering that Verovio also allows for variants encoded in MEI to be rendered natively and without having to be converted to another format.

The screenshot displays the MODE project interface for a madrigal. The title is "4. Anima cruda sì, ma però bella" by Battista Guarini. The interface includes a sidebar on the left with sections: "Display" (Zoom: 40%), "Page layout" (Horizontal/Vertical), "Options" (Modern clefs, Original clefs, Show critical notes), "Recording" (Luca Marenzio IL PASTOR FIDO), and "Tools" (PDF). The main area shows the musical score for five voices: Canto, Alto, Quinto, Tenore, and Basso. The lyrics are in Italian, and a recording player is visible at the bottom of the score.

Figure 2. A screenshot of a madrigal as displayed in the MODE project. The music is rendered dynamically with Verovio. For this particular madrigal, the music notation is linked to a recording of the piece that can be played directly in the interface.

3.3 The development of Git and related tools

Hosting a digital edition project is not an easy task to organise. Research grants usually do not provide resources for hosting beyond the duration of the project itself. This often means that hosting is organised at an institution where a project member works. However, it is not rare for scholars to move away from an

institution, and hosting becomes again a problem to be solved. Sometimes, other changes can occur at institutions, and some hosting solutions become no longer possible.

The MODE project had to re-organise its hosting several times. Because of the large quantity of images and files the project had to deal with, it was necessary to have an appropriate data hosting solution. This was provided by the University of Geneva and then by the Distributed Digital Music Archives & Libraries Lab (DDMAL) at McGill University.¹⁵ Eventually, all of the hosting was moved to the Center for Digital Research and Scholarship (CDRS) at Columbia University.¹⁶ Hosting of the images there was provided via Alfresco, an enterprise-level system for managing large collections of data. This allowed for all the digitised microfilms the MODE project had to deal with to be stored in one place and to be easily accessible by all members of the project. However, in 2015, the CDRS ended the support of the Alfresco system and all the data was moved to GoogleDrive.

The hosting of the project website also went through multiple institutions. At the beginning of the project, the needs were limited to a simple website. However, with the shift to produce a digital edition and the progress of the project, a more elaborated solution had to be put in place in order to publish the digital edition itself. As of 2006, Harvard University hosted the project webpages, which were then moved to the DDMAL at McGill University and then to the CDRS at Columbia University in 2014. The solution adopted at the CDRS for hosting both the website and the digital edition was the WordPress content management system (CMS).¹⁷ The WordPress CMS was used for the publication of the output of the NEH project with a pre-release. At this stage, the music was still being published as PDF files generated from the Sibelius files and not through Verovio. In 2016, however, the CDRS ended the support of the WordPress CMS and another solution had to be found. It was decided with the staff at the University of Columbia to move the entire content of the edition to a Git repository.

Git is an open-source software for managing versions of documents.¹⁸ It distinguishes itself from other widely used similar tools such as the Apache Subversion¹⁹ by being decentralised. This means that it does not rely on a centralised server for managing external contributions but instead simply allows for fully independent repositories to communicate with each other. This makes Git repositories much simpler to archive or to move, since each repository includes its entire

15 DDMAL, <https://ddmal.music.mcgill.ca>.

16 CDRS, <http://cdrs.columbia.edu>.

17 WordPress, <https://wordpress.com>.

18 Git, <https://git-scm.com>.

19 Apache Subversion, <https://subversion.apache.org>.

history and can interact with others in a completely independent and flexible way.

The documents typically managed in Git repositories are computer source code files. The use of Git, however, is not limited to computer source code files. Git is in fact perfectly suitable for text data in general, and this of course includes humanities data. The same applies to many functionalities offered by software that was built on Git. Some of these functionalities are directly optimised for managing computer source code, such as specific computer languages syntax highlighting. Many of them, however, are not that specific and can also be applied and be useful to other types of files. For example, the ability to track changes and to visualise all the revisions made to a file is very useful in many cases outside computer programming workflows. This also applies when it is necessary to merge changes made to a same file simultaneously by different contributors. In other words, digital music edition projects, and digital humanities projects in general, can benefit here from a side effect and make direct use of tools that were not meant to be used in this context.

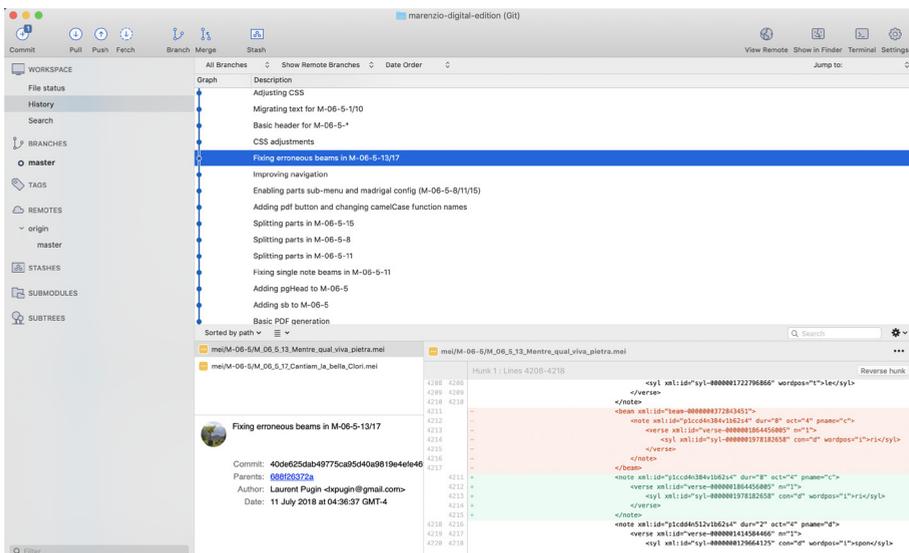


Figure 3. A Git repository browsed through a dedicated software application. Changes and revisions are all listed and differences between versions can be browsed. Additions and deletions are listed and highlighted.

For Git, this goes one step further. The development of Git is closely related to that of GitHub, a service for hosting Git repository.²⁰ The service is used by thousands of companies for managing their codebase, including major technology companies. For open-source projects, this service is free, with the implicit requirement that the source code (e.g., the files) are made accessible to everyone.²¹ For projects such as MODE which have open-data and open-source development policies it means that data hosting is available free of charge through the indirect support of private companies that use the service through paid plans.

In the case of the MODE project, GitHub appeared to be even more appropriate, since the service includes a web-publishing framework that perfectly fills the needs of the project. This web-publishing platform, named Jekyll, acts as a static content management system (CMS) fully embedded in the Git repository.²² The MODE project uses it to publish both its website and the complete digital edition – currently at a beta stage. With this approach, one single Git repository contains the complete MODE project, i.e. all the MEI source files of the edition content and also the complete tool ecosystem used to publish the edition. It makes the whole infrastructure very flexible, also for moving the hosting to another provider should this be necessary in the future. The repository also includes the entire history of all the revisions made to all the files. This allows for a previous stage of any files to be consulted or restored easily at any time.

4. Use of Aruspix for the editorial work

One of the main components of the NEH project was to use Aruspix for the preparation of the editions and more precisely to compare the original sources in order to elaborate the critical apparatus. However, the core of the NEH project was to support the editorial work and it did not support any software developments as these were explicitly precluded by the »Scholarly Editions and Translations« program. It did support, however, some training workshops, which were very important for the users to learn how to use Aruspix.

The use of Aruspix was two-fold. Firstly, it was used for comparing different copies of an edition. In Aruspix, copies of the same edition are compared through image superimposition. Images are automatically aligned using image processing

20 GitHub, <https://github.com>. GitHub is the largest provider for Git repositories but not the only one. It was founded in 2008 and currently hosts 80 million repositories. In 2018, GitHub was acquired by Microsoft for \$7.5 billion USD. See <https://en.wikipedia.org/wiki/GitHub>.

21 There is no requirement on the license under which the files have to be made available except for the fact that it needs to be compliant with open-source or open-data practices.

22 Jekyll, <https://jekyllrb.com>.

Figure 4. An example of the use of GitHub within its editorial workflow. It enabled the development and deployment of a specific online tool for editors of the MODE project to correct MEI files displayed interactively through Verovio.

algorithms specifically developed for this task. The superimposed images highlight differences between two images (e.g., two copies) in colour, making it easy for the editor to spot differences. This feature is very useful to detect stop-press corrections or other modifications that might have occurred during the printing process. For the Marenzio project, it was decided that this comparison would be performed only for the first edition (*editio princeps*) of each madrigal or canzonetta book, since this is the »copy-text« with which the successive reprints would be compared.²³ This feature of Aruspix was used to gather information about the printing process and to improve the edition of the music and the text by detecting mistakes. Clearly, the editor's critical eye is always involved in checking any outcome.

Secondly, Aruspix was used to compare the first edition with successive reprints. For this step, image superimposition is not a possible approach because the layout of the reprints is different. It can be quite similar in some cases, especially when a reprint was re-typeset in a mechanical manner using the *editio princeps*

²³ In some publications related to the project, the term »re-edition« is synonymously used for »re-print«. With moveable types, a reprint does require the content to be re-typeset, which in many cases involves corrections or modifications. For this reason, we can argue that both terms are equally appropriate.

as a model but never perfectly identical and image superimposition cannot work. For a digital comparison to be possible, the content of the prints that need to be compared first has to be recognised. Only then can the recognised content be digitally collated. Aruspix includes workspaces to perform both these tasks, namely an OMR workspace for transcribing the sources and a digital collation workspace for comparing the digitally transcribed sources. As for the image superimposition comparison technique, a visual check by the editor is needed throughout this phase too. The results are then used by the editor to write the critical apparatus.

The result of the comparison of all remaining copies of the *édition princeps* and the comparison of one copy of all known reprints with the *édition princeps* was gathered by the editors in the critical apparatus. It includes critical notes for the music which describe the discrepancies highlighted by the comparison through Aruspix and critical notes for text differences. In addition to listing all the differences in the various sources, the goal of this extensive source comparison approach was also to gather a better understanding of the editorial and printing process, not necessarily for Marenzio's work specifically but for music at that time in general. One take on this is that looking at differences and small anomalies can sometimes be the best way to learn about how normal workflows or normal tasks were organised.²⁴

What lessons can we learn from the use of Aruspix in the Marenzio project? How did all this work? First of all, the comparison presented above underlined how important it is to have easy-to-use tools. Indeed, one of the strengths of the Aruspix software application is to have a user-friendly interface for running the processes and correcting the data. The user interface of Aruspix implements common software application user interface design patterns, which makes it easy-to-use for editors. A significant amount of development time was put into this in the development of Aruspix, and it paid off when using it on this large-scale and practical project. Developing user-friendly and well-documented software application user interfaces is a well-known challenge for tools developed as part of research projects, such as Aruspix. Developing this is often difficult to include in a research project, mostly because it usually does not belong to the core goal of the research. It is sometimes simply not seen as a research component. The fact that Aruspix could be used successfully in the MODE project shows a real accomplishment in this regard.

24 This is not dissimilar to approaches in pathology studies in psychology. A pioneer for this practice was Théodule Ribot in the nineteenth century with his monography *Les Maladies de la Mémoire* published in 1881. Ribot's goal was to unlock the secret of the structure and functioning of normal memory through the study of its pathology. The pathological, by the excesses or the lacks that it presents, makes it possible to understand the normal functioning of the organism.

Nonetheless, the Aruspix software application remains not straightforward to use without dedicated training. The tasks to be accomplished with it remain uncommon and specialised (e.g., building a digital collation of OMR output is not a common task that users are used to go through). The training workshops planned and held during the NEH grant were absolutely essential for editors to familiarise themselves with the software and with the workflows. They were also important for setting up rules and schemes on how to organise and name the hundreds of files that the process uses (input image files) and generates (recognition files, training data files, collation files, etc.).²⁵

Another lesson learned from the project is that it is necessary to expect to have sources that will not be processable. This can be due to a too severe degradation of the document itself, since Early music sources often present deep bleed-through, i.e. when over time the content of the other side of a page becomes visible through the paper.²⁶ In some cases, the degradation is so far progressed that Aruspix is not able to recognise the content despite the extensive research that has been made to tackle this issue.²⁷ In other cases, the process can fail when the quality of the scanning does not meet the desired standard. This can occur with images scanned from degraded microfilms or when the scanning settings were not appropriate, meaning a very low resolution, a black and white (e.g., non-greyscale) colour scheme, or even both. The most problematic cases, however, are the images of sources not digitised with a proper scanning infrastructure but instead simply with a digital camera. When the images are taken without a tripod and the shooting angle is not perpendicular, the page in the image will be distorted perspectively. A small degree of distortion is tolerable in the process, but a too strong distortion will cause the staff detection algorithm in Aruspix to fail. Images taken without a proper scanning infrastructure often present problems related to the lighting. It can be non-uniform or simply too low, resulting in images lacking contrast. In other cases, not using a tripod can cause the image to be blurred because the person moved when taking the picture. These are cases we encountered in the images we received from some libraries in the MODE project.

25 The fact that Aruspix is a desktop application required it to be installed on all the editor's computers, which was not always easy to achieve. However, this never held up the project. With online web applications, which are more and more becoming the architecture used for distributing research tools of this sort, this would have been less of an issue. However, workshops and training sessions are going to remain necessary for any similar project to be successful.

26 With music documents, bleed-through is often stronger than with text documents because some symbols in music notation, such as filled note heads, cover a large area on the paper.

27 John Ashley Burgoyne, Laurent Pugin, Greg Eustace, and Ichiro Fujinaga, »A comparative survey of image binarisation algorithms for optical recognition on degraded musical sources,« in *Proceedings of the 8th International Society on Music Information Retrieval Conference (ISMIR 2007)*, Vienna, Austria (2007), pp. 509–12.

However, only a couple of books could not be compared automatically with Aruspix and had to be compared visually in a traditional manner.

The overall experience of the project is that the source comparison by OMR and digital collation remains a complex workflow to put in place. In addition to the complexity of the workflow itself and the problem that the source degradation or the image quality cannot be guaranteed in advance, there is also the fact that, as it stands, building digital collation in Aruspix still requires quite some manual work for specifying which page in an edition has to be compared with which other page of the reprint. This is quite straightforward when each piece is always printed on one single page, and in such cases the amount of manual work remains very limited and acceptable. However, this is not always the case, and preparing the digital collation can become a quite significant task in the process. Another important limitation with the current developmental stage of Aruspix is that the lyrics are not taken into account in the digital collation process. This remains left to the editor and needs to be done visually. This limitation is particularly problematic with madrigal repertoire simply because variants or corrections in the lyrics of this repertoire are noteworthy. They are certainly as important as variants or corrections in the music, if not even more so.

On the one side, the experience in preparing MODE showed that comparing reprints using OMR and digital collation at the scale of a full editorial project remains a significant task for the editors despite the very good results yielded by the Aruspix software for recognising the music itself. Considering the limited findings in terms of variants or corrections brought to light by the process, one can reasonably argue that at the current developmental stage of the technology the endeavour remains very expensive, considering the required worktime of such a project. This of course does not invalidate the approach itself, nor does it set aside the hope to see in the future improved workflows allowing for sources to be exhaustively compared automatically at a more reasonable worktime cost.

On the other side, however, the project confirmed that comparing sources by superimposition has to be more than highly recommended. This is fairly easily done once the images are available. Furthermore, the degradation of the sources has a lower impact on the workflow. The superimposition algorithms are much simpler and much more robust to document degradation such as bleed-through. The process does not involve the systematic verification of the OMR results before building the collation as it is the case with the comparison of the reprints. Finally, the pages can be superimposed independently of their content, i.e. without having to specify manually where a piece starts or where it ends because this does not affect the image superimposition process. Overall, the process for comparing copies of an edition is quite straightforward in comparison to the work-

flow for comparing reprints. Its usefulness is reinforced by the fact that, when applied to the first edition, it is also much more likely to bring to light important differences between the copies than differences between subsequent reprints. The limited findings made when comparing reprints in MODE confirm that reprints were most often copied from a previous version in a quite mechanical manner and without significant editorial revisions. For the first editions, on the contrary, the comparison of the copies confirmed that intervention within the printing process was quite common. Finding all stop-press corrections in the first edition should not be avoided by editorial projects aiming at producing high-quality critical editions. Stop-press correction can be of various kinds and in some cases can be the key to understanding what happened during the original editorial process.²⁸

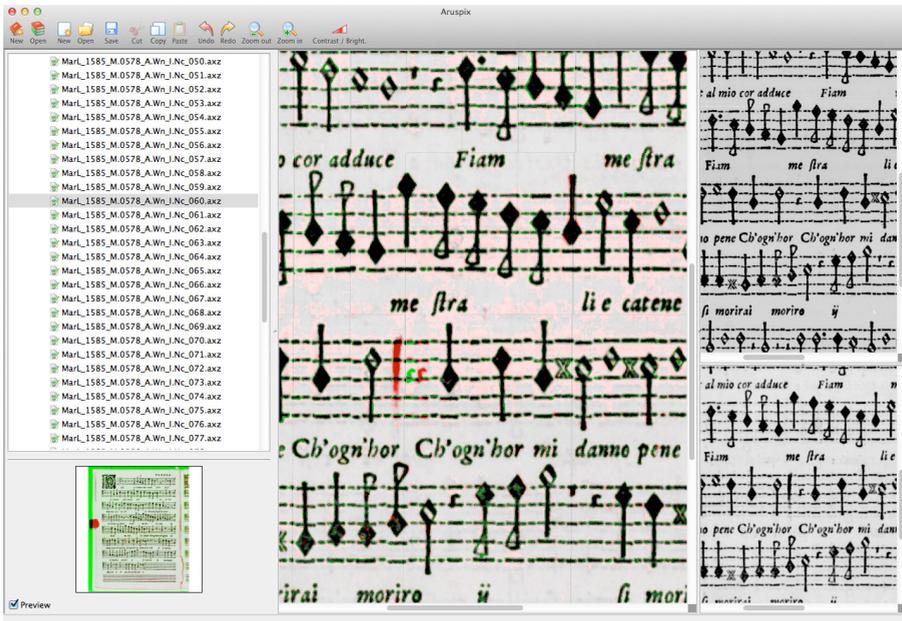


Figure 5. The comparison of different copies of the *edition princeps* revealed itself to be the most valuable operation in the editorial process of the MODE project. The automatic superimposition algorithm performs quite well even with fairly degraded sources and no additional work is required for the human editor to look for the differences that are highlighted in colour.

²⁸ For an example of the differences identified in MODE, see Laurent Pugin, »Music Printers at Work« (cf. fn. 4)

5. Leveraging the full potential of OMR transcription

The output of the part-books transcribed with Aruspix for the comparison of the sources remains monophonic. The logical next step in an editorial process would be to use transcriptions resulting from the OMR output to directly build the score. This was planned in the MODE project as a follow-up of the NEH grant. However, due to the lack of funding, it was not possible to put this into place. What remains complex in this process? Which steps do we need to go through for moving from the OMR output to the score and will this be possible in the near future? While this looks conceptually like a straightforward task where parts simply need to be put together, different parameters listed below still make it quite challenging.

The first challenge for assembling the parts into a score is to determine which pages belong to which pieces. In some cases, when the part-books are structured in a very systematic way, the task can be quite straightforward. The part-books were frequently organised systematically because the music printers used to favour it. For example, in order to make the printing process as streamlined as possible, printers always tried to apply the same layout pattern. This made the typesetting process simpler and even allowed for keeping text elements on the printing form from one page or one part to another. This can be observed easily from typesetting errors such as wrong page numbering or any other anomaly such as broken types that often appear repeatedly on all pages or on all parts in exactly the same place. The best-case scenario for the printer, and for us today when processing the books, was to have one piece (i.e., one madrigal) per page and also to have the same number of madrigals throughout the books. Each part-book had the same number of pages (or folios) with the same page numbering and also the same table of contents, which they did not have to re-typeset for each part-book. The same form could then simply be reused for all part-books.

The situation quickly becomes more complicated when such a systematic organisation of the part-books is missing. It also often happens that each part of a piece is not printed on a single page but across several pages. Conversely, one page may contain more than one piece or part of a piece. Typically, with madrigals of various lengths, printers would try to optimise the use of the paper and fill the pages as much as possible. A long madrigal can be printed on more than one page, with a staff or more on the next page, followed by a shorter one that would fill the rest.

In the future, it would be interesting to integrate an algorithm in the workflow that can detect the beginning and the end of pieces throughout a part-book. It would then be possible to use the output of these algorithms in order to prepare

the assembling of the parts without having to individually label the manual annotations, or at least with significantly less work.

Another complication arises when the number of voices in the pieces is not consistent throughout the book. The *Sesto libro de madrigali a cinque voci* (Venice: Angelo Gardano, 1594; RISM M-557) exemplifies this, with all madrigals composed for five voices except for the last, *Cantiam la bella Clori*, which is set to eight voices in two choirs of four voices each. Because there are only five part-books, the madrigal at eight voices is printed as what can be interpreted as a mixed mode of part-book and choir book, with two voices in some of the part-books. The *Canto* part-book contains *Canto* and *Tenor* of the *Primo choro*, the *Alto* part-book the *Alto* and *Basso* of the *Primo choro*, the *Tenor* part-book only one voice, namely the *Tenor* of the *Secondo choro*, the *Basso* only the *Basso* of the *Secondo choro* (indication is missing), and the *Quinto* the *Canto* and *Alto* of the *Secondo choro*. Where more than one voice is contained in a part-book, two singers are expected to share a part-book in the same way singers share a choir book. Although this represents a quite common practice, these types of irregularities do make automatic assembling of parts from OMR output quite challenging.²⁹

Finally, with Renaissance music in general, assembling part-books is made even more difficult when the music is in ternary meter. In Marenzio's secular music, however, this is not a significant challenge because sections in ternary meter are limited and do not include complex proportions. The number of alterations or imperfections is very limited and can easily be handled by hand. Nonetheless, recent work in the field showed that resolving mensural rules algorithmically is possible and that it is perfectly feasible to integrate these algorithms into the score assembling workflow.³⁰

29 More generally, we can observe patterns in the layout of the printed books that made it easier for the printers to organise the printing process, from typesetting to binding. However, the typesetting layout of music notation was unconstrained, and the printers had the freedom to lay out the music as needed. In the simplest cases, with one piece per page and one page per piece, with the same number of voices throughout the book, grouping the parts to build the score is fairly trivial. With more complex cases, it is possible to envisage approaches using intelligent text-mining based on text elements of the part-books extracted via OCR, such as voice labels (*Canto*, *Alto*, etc.), lyrics, or piece titles. However, any of these solutions will certainly require human supervision and verification.

30 Martha Thomae Elias, »Automatic scoring up of mensural music using perfect mensurations, 1300- 1550« (MA dissertation, McGill University, Montreal, 2017).

6. Conclusion

The evolution of the MODE project illustrates well the need for digital edition projects in order to be able to react to technological advances continuously. When the project started, MEI, Verovio, and Git did not yet exist – although MEI already existed it was only as a one-person research project with no released specifications and no community support. Today, they provide the core of the MODE project. These innovations had a profound impact on the overall goal of the project. Throughout the years, by integrating the possibilities offered by these new resources, the project moved its target from a digital edition designed as a publication to an edition designed as an extensible open dataset and reusable open-source environment, i.e. extending the outcome of the project beyond just the edition itself by making all the sources of the edition, data, and tools available under a form that can be re-used and extended by other projects.

A key approach behind this new concept is the clear separation between data and tools. For a project such as MODE to be maintainable in the long run, this separation is absolutely essential. It guarantees that the data remain readable and usable even if one day the tools on which the edition is built are no longer maintained or no longer work. This separation is one of the many contributions brought by MEI and the continuing work of its community. For the MODE project, Verovio has opened up brand new possibilities in the way the music notation is made accessible to the reader. The modular approach of Verovio makes it possible for it to be used in a very wide range of environments and not only digital edition projects. By growing the community of users and also of developers who are willing to contribute to the code and its documentation, the project maximises its sustainability perspectives. Enabling contributions from third-party projects is a unique feature of open-source tools, and the more widely they are used, the more likely they will be maintained and supported by the community. Finally, with the use of Git and hosting on GitHub, the MODE project benefits from cutting edge tools for managing digital workflow but also from a reliable and flexible hosting solution made available for free by a very large technology company willing to support open-source projects. The publication tools integrated in GitHub offer a simple, static, and self-contained solution.

The experience regarding the use of Aruspix gathered throughout the NEH project for MODE is very valuable. On the one side, it highlighted that the comparison of reprints remains quite demanding in terms of organisation and human resources to be involved in the process. As discussed in the paper, developing algorithms for automatically detecting and processing the structure of the part-books (e.g., detecting pieces, detecting voices) will significantly reduce the

resources needed for comparing reprints. Furthermore, these algorithms will also help to facilitate the use of the OMR transcription for generating scores from them even though we have seen that this process requires to shift the viewpoint at various steps. This is not always straightforward. Automatically assembling scores will undoubtedly be useful for large-scale approaches, a domain where the use of Aruspix still has considerable potential, and various research projects outside the digital edition realm already benefit from it.³¹ On the other side, the use of Aruspix in the MODE project confirmed the efficiency and the usefulness of comparing existing copies of the first edition in particular. Not only is this comparison relatively easy and not highly time consuming, but it also promises to bring to light significant and relevant differences.

The MODE project will have to continue to evolve and to adjust itself to innovations. What will they be? This is of course difficult to predict. However, changes in various domains, and not exclusively scholarly ones, have shown us in the past that we tend to overestimate the changes that will occur in the next five years, while we simultaneously underestimate the ones that will occur in the next ten years. It sounds indeed fairly likely that the resources on which the MODE project is built will still be at its core in five years. It is not unlikely that this will also be the case in ten years even though this is clearly more difficult to predict. For MEI, we can see shifts occurring from the traditional quite monolithic modelling approach to solutions making more use of Linked Data (LD) technologies. These will quite undoubtedly be a key technology that will enable the structuring of data and the related tools to be rethought and adjusted. In this context, the conceptual shift the MODE project went through, moving from an edition designed as a publication to a digital environment where the dataset and the publication framework can continuously evolve, appears to be appropriate and will hopefully offer the required flexibility.

31 See the Single Interface for Music Score Searching and Analysis (SIMSSA) <https://simssa.ca/> and the Full-Text search of Early Music Prints Online (F-Tempo) <https://f-tempo.org> projects.