TROMPA: Towards Richer Online Music Public-domain Archives

Deliverable D2.1

Early Requirements

Grant Agreement nr	770376
Project runtime	May 2018 - April 2021
Document Reference	TR-D2.1-Early Requirements
Work Package	WP2 – Requirements and User-Centered Design
Deliverable Type	Report
Dissemination Level	CO-Consortium
Document due date	31.10.2018
Date of submission	31.10.2018
Leader	PN
Contact Person	Vladimir Viro (vladimir@viro.name)
Authors	Vladimir Viro (PN), Natalia Barsova (PN), Tim Crawford (GOLD), David Weigl (MDW), Cynthia Liem (TUD), Nicolás Felipe Gutiérrez Páez (UPF), Patricia Santos (UPF), Álvaro Sarasúa (VL)
Reviewers	Emilia Gomez (UPF), Aggelos Gkiokas (UPF)

Executive Summary

The Trompa project strives to create systems for multiple audiences, involving these audiences during the planning, soliciting their feedback during the iterative refinement and feature delivery process, as well as inviting them to contribute their musical knowledge using the created systems. The early requirements planning concerns itself with the first part – soliciting user feedback for the first version of use cases developed by TROMPA. This deliverable comprises the requirements elicitation methodology, the resulting application scenarios and user stories and high-level functional/non-functional requirements that are based on application scenarios and user stories developed in the tasks T2.1 and T2.2. Deliverables D5.1 and D4.2 draw upon the functional requirements and application scenarios covered in this deliverable, as well as, naturally, most other following deliverables.

We focus on five user communities: music scholars, music performers, orchestras, choirs, and music enthusiasts.

The requirements and feedback elicitation from our target user communities has started before the project kick-off. Through a collaboration between the TROMPA partner PN and the project's associated partner IMSLP, the largest online sheet music library, we have conducted a survey among thousands of IMSLP users. People with music backgrounds ranging from music enthusiasts, to music students, teachers, professional musicians, conductors, librarians, and music researchers, have contributed structured and unstructured feedback to questions about the desirability of certain functionality that we have proposed to develop for them. This feedback (summarized in Appendices B and C of the present document) has influenced our planing, and the survey participants who expressed interested in being informed about further project developments (over eight thousand people) will be tapped for more feedback and invited to be our beta testers as the development of our use cases progresses. Section 2 of this document addresses requirements elicitation in more detail.

In Section 3 we describe our user communities in more detail and state the preliminary requirements for the use cases we plan to implement for these communities.

In the case of music scholars, our main focus lies on sheet music and the features around annotating the scores, linking the scores, searching for content either by text or musical content itself (for example, selecting a section of a score to find places in the current or other scores containing the same melody or chord progression), etc.. We also provide facilities for digital enquiries by making the search API publicly available and publishing the bulk of data collected in TROMPA as a public linked open data dataset.

In case of music performers we are focusing on piano players, leveraging technology by PN and community access of MDW to start early beta testing of the system with piano students at MDW already in the first months of 2019. We intend to record their complete practice sessions and provide them feedback on the way they practice. We also let them access to the CDR recording library in a novel way – they can easily compare the way they interpret a piece with the interpretations from the CDR archive, which holds a very large collection of piano music. We provide performers functionality for annotating scores with comments and practice suggestions, which can be publicly shared.

For the choir use case we develop functionality both for choir conductors and choir members. Choir conductors will be able to create repertoires importing scores in MusicXML or MEI format and to share these repertoires with choir members. Choir members will be able to practice at home loading any of the scores in the repertoire. For practicing, they will get piano-roll and score visualizations of the different voices in the piece and they will be able to choose one of the voices. Then, they will be able to get synthesized versions for the rest of the voices and to sing-along with them, controlling their respective volumes and panning. The voice synthesis will support multiple languages. After practicing, choir members will get automatic feedback about their performance for different aspects (timing, intonation, etc). Choir conductors will be able to retrieve and listen to performances by choir members for a particular score in the repertoire, and they will be able to provide choir members with feedback on their performance.

For the orchestra use case we are focusing on creating high quality open digital sheet music editions of out-of-copyright music using a crowd-sourcing approach. Digital score editions that allow for the addition of digital annotations will enable richer archiving of performance notes together with the scores they were based on.

To music enthusiasts we provide a new way of exploring music recording libraries by visually comparing various interpretations. This comes in addition to the score annotation functionality that allows them to discover content prepared by music scholars, musicians and other enthusiasts. Likewise, we provide a novel learning community in the form of a MOOC that aims to engage music enthusiasts in expanding their knowledge about European classical music heritage. The course will also introduce how novel technologies, as the ones proposed by TROMPA, are used to analyze music resources and enhance the current status of music reposities by involving citizens in crowdsourcing tasks. The MOOC will support the connection and collaboration with other music enthusiasts through specific activities encouraging active participation.

It should be noted that, since the user communities overlap (for example, performers may very well be interested in functionality developed for music enthusiasts or music scholars), the use requirements defined for one community may be relevant to other communities and use cases as well. After listing the functional requirements shared by all use cases (such as data export, account system, and internationalization) and listing functional requirements specific to each use case, we provide an overview of the overlapping functionality using a matrix with rows corresponding to functional requirements and columns corresponding to our use cases.

The use cases share some common non-functional requirements, such as GDPR compliance and open access to public data generated by the use cases. Use case specific non-functional requirements include data processing performance necessary to ensure timely feedback for performers, content quality guaranteeing the usefulness of the created digital scores for orchestras, as well as the printability of the digital scores with annotations, and the ability of performers to use the system on mobile devices when they practice away from computer.

Version Log		
#	Date	Description
v0.1	09.10.2018	Initial version submitted for internal review
v0.2	19.10.2018	Revised version after initial comments
v0.3	31.10.2018	Revised version after internal review
v1.0	31.10.2018	Final version submitted to EU

Table of Contents

1. Introduction

The general goal of the TROMPA project is to broaden and deepen the enrichment of online public-domain classical music resources. This is done by engaging a broad audience of music-loving citizens, and having them strategically cooperate with novel ICT technologies.

As specified in the TROMPA Description of Action, five main target audience categories are of particular interest to TROMPA: music scholars, orchestras (content owners), instrumentalists, choir singers, and music enthusiasts. Part of the project (as researched under WP 2, 5 and 6) will focus on use cases relevant to these audiences. This involves identifying audience-specific requirements, recognizing and acting upon any common desired elements in terms of data representation and infrastructure, and ultimately proposing several audience-specific pilots. Within WP3 and 4, intelligent music processing technology and crowd engagement mechanisms will be researched, that will be integrated in the TROMPA workflow to improve scalability of musical enrichments.

This deliverable reports on early requirements, as established by consortium partners and representative users in the first months of the TROMPA project, and corresponding to tasks T2.1 and T2.2 in the TROMPA Description of Action. This involves an initial requirements elicitation phase (described in section 2), a first definition of relevant use case scenarios per audience categories, including user stories (section 3), and an initial listing of high-level functional/non-functional requirements (sections 4 and 5).

The listings and stories should not yet be seen as complete and finalized; instead, they follow a first basis for further refinement with target audience and partners. A full and final requirements document will follow in M18 of the project, as D2.2 Final Requirements.

2. Requirements Elicitation

Before the TROMPA project started, partner PN had already initiated requirement elicitation as part of the Music Connection Machine survey, targeted at users of the Petrucci Music Library (IMSLP). With IMSLP being a widely used music score resource, the Music Connection Machine (MCM) survey collected more than 11600 survey responses from music scholars, musicians, students, teachers and music enthusiasts from around the world, in which participants emphasized their interests in particular developments that fall in TROMPA's scope and provided free-form feedback.

In this section, we give a first analysis of responses and common topics in the answers to this survey. We also describe the strategies to keep engaging the MCM audience, and have their feedback connected to requirements as obtained through TROMPA's partners and audience stakeholders.

2.1 Music Connection Machine survey

We have promoted the MCM survey, available at http://www.musicconnectionmachine.org/, on the website of IMSLP, the Petrucci Music Library, through an ad in the Peachnote Score Viewer. The survey has been online for over three years and has collected more than 11,600 responses. The survey has asked the respondents whether several specific functionalities would be interesting to them, and also provided ways of leaving free form feedback, which has been submitted by about one quarter of all respondents. Over 8 thousand respondents have provided an email address to be contacted later, once the project proposed by the MCM, has started. We have also learned about the demographics and the target user group makeup of our audience.

Responses to the six yes/no questions of the kind "Would you be interested in FeatureX?" (see Appendix A for the complete list of these questions) have provided us with confidence in the chosen project direction - for each of these questions the approval rate has been around 90%, with slight variation (see Appendix B). The free form responses have been a great source of additional ideas and insight about the user priorities. We have summarized and used them to guide our planning process (see Appendix C, which is referenced by the relevant functional requirements in section 4).

One should take into account the context, in which the survey has been promoted – the online sheet music library (imslp.org). The respondents tend to write more about issues surrounding music works and sheet music editions
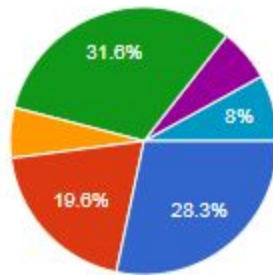
Figure 2.1 Distribution of the MCM survey participant self-identifications

2.2 First follow-up with survey respondents

Given its size and notable engagement, as well as the clear overlap between interests of the MCM survey respondents with those of TROMPA, it will be valuable to engage the MCM audience for TROMPA.

Because of the GDPR, the MCM mailing list (containing email addresses of over 8200 recipients) needed membership reconfirmation. As the GDPR came into effect a few weeks after the TROMPA kick-off meeting in Barcelona, partner PN prepared and sent out a message to the MCM mailing list, inviting subscribers to re-subscribe, while also informing them about the TROMPA project, and announcing an upcoming follow-up requirements elicitation survey.

2.3 Follow-up survey strategy

Next to the input from the MCM survey, requirements for the TROMPA use cases are informed by input from consortium partners and use case audiences (which are either represented by consortium partners, or reachable through these partners).

At present, the TROMPA consortium partners defined representative user stories, functional and non-functional requirements for each of the use cases. These will form the basis for the planning of concrete enriched user-facing pilot applications. In moving towards such applications, surveys at a more detailed level than the MCM survey (in terms of dedicated audience targeting and requirements elicitation) will be formulated within the consortium, and circulated to the dedicated use case audience categories. These surveys will be disseminated through consortium members, associated partners, and the TROMPA communication channels as defined under WP7.

The collection of user feedback is an ongoing process, that will go beyond the early requirements collection stage. After every major development iteration, we will contact and update people from the corresponding user communities who consented to being contacted, as well as the beta testers for use cases, for which we have already distributed our preliminary development outcomes to selected users, soliciting their feedback on newly developed functionality.

2.4 Beta testing

Another way of engaging with our user communities as early and closely as possible is inviting people to participate in beta testing the use cases we are developing and learning from their usage patterns and their feedback, directly provided using the application or collected through surveys. The first such beta testing proccess will be started in the spring of 2019 at the MDW for the performers use case (cf. section 3.3), with further use cases making use of this approach later on as their technical prerequisites are implemented.

3. Application Scenarios and User Stories

In this section we describe the application scenarios for the five user communities of main interest to TROMPA: music scholars, orchestras, performers, choirs, and music enthusiasts. The communities are not disjoint, for one may count many performers as music enthusiasts as well, some performers may share interests with music scholars, choir singers with performers, etc.

3.1 Music scholars

Music scholars – i.e., those involved in professional or amateur pursuit of deep knowledge about music for its own sake rather than for commercial reasons – are interested in finding connections between music works on multiple levels: from co-occurrences of melodies, harmonic and rhythmic progressions, to the large-scale structural similarities of musical works. Sheet music is the main starting point and anchor for such research. Existing sheet music editions and differences among them can also be a subject of analysis.

We need to provide scholars with tools for linking between scanned sheet music, searching and navigating based on scanned sheet music, as well as digital sheet music encoded in the MEI format.

In order to discuss higher-level structure of music works, a timeline presentation also needs to be available to the users, and it should be possible to create and interact with segmentations of music works on this level.

Audio and video performance artifacts should be referable both on the sheet level and the timeline level. Users should be able to illustrate certain passages in the music using external audio and video material hosted on commercial on non-profit public platforms, such as YouTube, SoundCloud, Muziekweb, Wikipedia, or Archive.org, as well as images, links to external websites, and linked data entities. The referenced sheet music can be hosted anywhere on the web as linked open data. TROMPA should provide an implementation for the web browser environment.

The user representative for this use case is GOLD, the technical representative is VD.

Use scenarios

As a music scholar, *I want* to illustrate differences in performative interpretation of a particular score section by reference to audio examples *so that* I can evidence a point of scholarly discussion.

As a student of early music, *I want to* find other vocal versions of a certain French chanson or Italian madrigal and also their arrangements for instruments, and I need to be able to find recorded examples and scholarly writing about the music, *so that* I can more fully understand its historical context.

As a musicologist, *I want* to supply human reference annotations of many different performances of the same work to provide "ground truth" *so that* an MIR researcher can train MIR recommender algorithms.

3.2 Orchestras

While orchestras are content owners with respect to their performances, they usually are not with respect to the scores they play, even if these scores consider works that are part of European cultural heritage in the public domain. The countless copyright-free orchestral scores are not as easily accessible as they should be. While the Petrucci Library (imslp.org) offers an incredible amount or scores, these are only scanned PDFs, often of a quality not suitable for performance at a professional level.

Frequently, those who would like to invest in decent performing material are faced with publishers who offer only (poor quality) rental material, and charge extra for the right to record/broadcast. For original musical works that are in the public domain, this should not be a requirement.

From a more practical viewpoint, digital score editions that allow for the addition of digital annotations will enable richer archiving of performance notes together with the scores they were based on.

The user representative for this use case is RCO, the technical representative is TUD.

Use scenarios

As an orchestra, **we want** digital editions of public-domain musical works, **so that** we can generate high-quality, playable, performance-ready instrumental parts.

As an orchestra, **we want** to digitally archive performance-relevant annotations that were e.g. made by librarians and conductors, **so that** these annotations can also be available to future performers.

As an orchestra, **we want** to be able to compare different historical editions of a work used by our past conductors, **so that** we can better understand the performance-specific interpretations made by these conductors.

3.3 Performers

In a local performance context (on stage, or in a practice room), the multimodal score-interaction interface provides classical musicians with a "performance companion" capable of analysing musical and performative features (e.g., for pianists: onset timings, dynamics, key position trajectories) in real-time.

Through audio/score alignment, these features can be visualised and contextualised within a musical score.

By virtue of interconnections with other score-aligned performance recordings of the same piece, performances can be compared along these features (and aggregate metrics, e.g., tempo). This enables performers to compare their own performances to reference recordings; or, to track the evolution of their own performances over time.

Feature data and score alignments (as well as performance recordings, with the performer's permission) are stored for future comparison, as well as for re-use in other contexts: to provide teachers insights into their student's performative progress; to enable enthusiasts and scholars to annotate performances, and incorporate them as evidence within scholarly discourse; and to provide music information researchers with performance data for training and testing of MIR algorithms.

Performers can comment on certain places in the scores, for example with practice or fingering suggestions, and these comments can be shared and seen by other users.

We develop this use case focusing mainly on piano players, although many planned functionalities can and will be extended to other instruments (e.g., guitar) straightforwardly as the project progresses.

The user representative for this use case is MDW, the technical representative is PN.

Use scenarios

As a student pianist, *I want* to compare my practice recordings of the first movement of Beethoven's Piano Sonata No. 23 to a 1960 reference recording by Richter *so that* I can learn to emulate Richter's stylistic timings in my playing.

As a student pianist, *I want* to annotate my practice performances in relation to the score and to practice metadata *so that* I can more precisely discuss performative aspects with my teacher, or with peers (at University or in a MOOC).

As a teacher, *I want* to compare recordings of my student practicing the same piece over a period of several weeks *so that* I can visualise the evolution of their performance parameters at key passages, to provide better-informed guidance for their practice.

As a self-taught classical guitarist, *I want* to be able to find pieces of music to download that I can manage to play, *so that* I do not suffer the frustration and waste of time in struggling with music that is too difficult.

3.4 Choirs

The goal of the Choir Singers Pilot is to assist (amateur) choir singers during individual performance. It also provides functionality for the choir conductor to create repertoires and to listen to performances by choir members, providing feedback to them. Users of the pilot should be able to synthesize existing scores, to sing-along with the synthesized voices, and to receive feedback on their performance. The accompanying voices should be available for music in Spanish, Catalan, Latin, English or German.

The user representative for this use case is UPF, the technical representative is VL.

Use scenarios

As a choir conductor, *I want* to retrieve existing scores to create a repertoire I can share with my choir members *so that* they can practice their parts individually.

As a choir conductor, *I want* to listen to the individual performances of my choir members *so that* I can keep track of their work and provide them with useful feedback.

As a choir singer, *I want* to get realistic synthesized versions of my choir repertoire I can sing along with *so that* I can practice my parts individually without the need to wait to meet the rest of the choir.

As a choir singer, *I want* to get automatic feedback on my performance *so that* I can identify parts where I can improve.

3.5 Music Enthusiasts

Music enthusiasts are people without an extensive musical background but who are willing to engage with and learn about music. Based on interests and opportunities offered by consortium partners, two use cases are identified, that will both build upon research and functional component conducted across the TROMPA project, e.g. with relation to automated performance analysis.

3.5.1. Music recording libraries

Users of music recording libraries often have access to multiple recordings of the same music piece. However, there is no way for them to easily see the differences between the recordings: with an exception of recording duration, no other musical property of the recordings is accessible to the users through the libraries' user interfaces. With this use case we are providing users with a new interface to navigate, discover and compare recordings based on their musical content.

The user opens a page of an audio or video recording and sees an interactive visualization of other recordings of the same piece in the library, with every recording represented by a line chart of its local velocity relative to the given recording.

The user can move their mouse over the visualization and the metadata of the recording that is represented by the chart closest to the mouse pointer will be presented near the visualization. If the user wants to listen to this recording, a left mouse click starts the audio playback from the time corresponding to the x coordinate of the point on the plot. A double click causes the page of the other recording to be loaded.

The plot is scaled in a way that makes it possible to identify meaningful differences between most performances (e.g. using the standard deviation of the differences in velocities as a scaling factor).

The UI component providing this functionality also allows to select particular recordings for comparison with the given one, which allows the user to evaluate differences only between performances that are interesting to them, avoiding information overload for pieces with lots of recordings available.

Since the repository we are going to work with (by the TROMPA partner CDR) currently has Dutch as its first metadata language, with English, French and German translations planned to become available until 2020, we want to support CDR's Dutch-speaking users, as well as English-speaking users worldwide. Having a multi-lingual system and metadata available in further languages, we will work to serve French- and German-speaking users metadata in their language.

The user representative for this use case is CDR, the technical representative is PN.

Use scenarios

As a music enthusiast, *I want* to easily have an overview over different interpretations of a piece that interests me *so that* I can listen to recordings that are different from the interpretation that I'm already familiar with.

As a music enthusiast, *I want* to easily identify places in music pieces that are performed in many different ways *so that* I can learn more about musical practice.

As a music enthusiast, *I want* compare two or more particular interpretations *so that* I can learn more about differences between the way artists I'm interested in interpret the same piece of music.

As a music enthusiast, *I want* to communicate my opinions and preferences about recordings of my favourite musical works *so that* I can identify my favourites for later retrieval, discover new suggested recordings based on my expressed preferences, participate in community discourse, and broadcast my expertise and enthusiasm ("music as a social badge").

3.5.2. MOOC use case: 'Learning about European music heritage - Introduction to new technologies for exploring and enriching classical music repositories'

The goal of the MOOC for Music Enthusiasts is to provide novel and playful interaction mechanisms for musical cultural heritage content aimed at people without formal musical knowledge, but with interest in learning more about European music heritage. The learning content and related activities will be designed by TROMPA domain experts. The MOOC will introduce concepts focused on how technologies (mainly tools developed by TROMPA partners) can be used to explore and enrich classical music repositories. Based on the concepts of citizen science and inquiry learning, this use case aims to create a learning community around the technologies designed within the context of TROMPA for exploring and enriching classical music resources. Participants of this learning community will learn how these technologies can support novel analysis of classical music resources, so they will be able to use them by their own and in different and wider scenarios.

The interactions between users and the resources (tools, musical works, research questions, annotations) will be mediated through an hybrid MOOC in which participants will initially acquire basic knowledge and skills related to music and the developed technologies (following a traditional design of a MOOC or also known as xMOOC). Participants will be also asked to participate in a collection of connectivist learning tasks (known as cMOOC activities). These activities will propose some learning challenges (e.g. proposed by TROMPA domain experts) and participants will have to use the TROMPA tools to create new links (annotations) between the resources in repositories. An illustrative example of activity can be the described as follows: MOOC participants will be asked to use the music recording libraries to perform comparison of performances (see previous use case for Music Enthusiasts in section 3.5.1), and to share and evaluate other users analysis (quantitative values as a likert scale will be combined with qualitative comments to support discussion and debate among participants), in order to accomplish a specific learning objectives within the MOOC defined by TROMPA domain experts.

Participants in this pilot will be trained in the use of different music processing tools, in order to perform specific tasks to solve questions and challenges proposed by advanced users (e.g. TROMPA domain experts or participants with advanced competences) related with the enrichment of specific musical works. They should likewise be able to explore the annotations and links provided by other users to help solving a specific question or challenge, as well as to provide their own responses (annotations, links or ratings of existent responses). Furthermore, more experienced users should be able to define new questions and a desired response structure (annotation or link between different musical works, performances or sections of scores).

As any other cMOOC, this MOOC is built around the idea of connectivity among a group of peer-learners. Knowledge creation/distribution is decentralized, therefore at the end of the MOOC each participant should be able to propose new challenges related to the use of TROMPA tools. While creating this decentralized knowledge, users (MOOC participants) will provide annotations (opinions, ratings, comments) that will serve as valuable data for researchers, e.g. it can serve as ground truth data

for recommendation systems, or will provide evaluation to automated processes used within the context of TROMPA. Additionally, participants' contributions will be rated by the community itself. Thus, feedback mechanisms as rankings, social badges and/or public recognition will be used to incentivize users to be actively involved in the cMOOC activities. The more active participants will be considered as advanced users, and will be engaged to create their own questions and challenges to gain more recognition within the learning community.

The user and the technical representative for this use case is UPF-TIDE.

Use scenarios

- **As a** music education researcher, **I want** to propose new questions such as "What is the relation between Mahler and klezmer music?" or "What kind of parallels/contrasts exist between Mahler and Beethoven's attitude to popular music?" **so that** music enthusiasts can explore musical works and find, comment and link them to help solving the question.
- **As a** music enthusiast, **I want** to learn about classical music and available tools to enrich the exploration and use of the existing resources, **so that** I can contribute actively to help in different crowdsourcing tasks according to my abilities and expertise.
- **As an** advanced MOOC participant, **I want** to obtain access to the annotations and feedback made by MOOC participants **so that** I can assess their performance and learning improvements during the course / **so that** I can create new questions based on the learning context.
- **As a** MOOC participant, **I want** to explore and rank annotations of other participants about specific questions **so that** I can compare and self-assess my own annotations to improve my knowledge of the specific topic and help to assess the best answer of the research questions.

4. Functional Requirements

Based on the application scenarios and user stories of the previous section, early functional and non-functional requirements can be extracted. Part of these can be identified as requirements at the general level, while others are specific to certain user groups.

In the current and following section, a listing of these early requirements is given. This listing will be used to identify common elements (which, in terms of infrastructure, will be supported by the TROMPA Contributor Environment, see e.g. D5.1), and to design and develop concrete user-facing pilot applications, which will be part of WP6.

As mentioned before, these listings are by no means exhaustive and final yet, and still will be refined throughout the first year of the TROMPA project.

In the Section 4.1, we present the general functional requirements of the platform that are shared by all use cases. Section 4.2 describes special requirements of every use case in detail. Section 4.3 summarizes all use-case specific requirements in a common table, and shows overlaps in requirements between the use cases.

4.1 General functionality

4.1.1 Accounts

Users should have accounts for TROMPA's platform if they are contributing data to at least one of the use cases. They should be able to create accounts using email or social media accounts. Viewing data made publicly available by TROMPA should not require any account.

A TROMPA account should provide access to all TROMPA resources and functionalities.

There should be a single place for signing up for a TROMPA account.

Users should have a public profile that will be associated with the content they make publicly accessible on the platform, if they create such content.

User data must be managed in accordance with GDPR. Users should be able to suspend and recover their accounts, or delete their accounts completely, leading to removal of all their PII from the system.

4.1.2 Search by text

Users should be able to use search for finding relevant music pieces and their scores.

The search should support different spellings of composer and music work names.

Metadata should be searchable in English.

Search requests in Dutch, German and French should be supported as well, subject to the availability of the metadata in these languages.

The search engine should consider the content's popularity and the user's favorites list when ranking search results.

4.1.3 Localization

The developed UI components should support internationalization. All components should be available in English and support further localization (e.g. to Dutch, German and French).

The component language should be explicitly configurable through the JavaScript API of the component (the way it is embedded and configured on a website), so that the website integrating the component can control its locale programmatically.

When metadata related to musical works or recordings is available in multiple languages (e.g. in case of Muziekweb), users should be able to make use of it when searching or browsing.

4.1.4 Score annotations

Users should be able to add textual, audio or video annotations to scores. (Cf. Appendix C 2.3)

Users should be able to group annotations in layers that can be shown or hidden at once. (Cf. Appendix C 2.3)

Users should be able to toggle and filter annotations made by different annotators.

4.2 User group specific functionality

4.2.1 Music scholars

Embeddable content: Material prepared by the users should be embeddable on third-party websites. It should be possible for scholars and teachers to use this in support of a scholarly argument or in a pedagogical context, for example a MOOC.

Content access control: Users should be able to work with the system in private or public, and be able to control the access to their documents and annotations (sets of connections made within a set of scores and/or recordings).

Search by text: Users should be able to search for content using text, with the search engine resolving the queries using named entities or annotations linked within the documents to be retrieved.

Search by musical content: Users should be able to search for content by non-textual content, such as musical contents of sections of sheet music selected by the user. Upon selecting a section of a music score, the users should be able to choose the search mode (e.g. harmonic, rhythmic, or other kind of similarity) and submit the search query.

The search results should be available on a separate page and the request should be identified by the page's URL, so that users can share the search results page with others.

RDF data export: The search request and results should be exportable as RDF by the user.

Public search API: The search functionality should be exposed as a public API.

4.2.2 Orchestras

Score part generation: Users should be able to generate instrumental parts from a full digital score.

Annotation propagation to score parts: User should be able to propagate annotations on a full score to instrumental parts, and vice versa. (Cf. Appendix C 2.3)

Score annotation layers: Users should be able to archive annotations as a layer with the corresponding digital score. (Cf. Appendix C 2.3)

Recordings as score annotations: Users should be able to archive a digital score together with their corresponding performance recordings. (Cf. Appendix C 2.4)

Score edition alignment: Users should be able to toggle edition variants of scores that correspond to the same underlying musical work. (Cf. Appendix C 2.2)

Filtering annotations by metadata: Users should be able to toggle and filter annotations made by different annotators.

4.2.3 Performers

Practice statistics: Users should be able to see statistics of their practice overlayed on top of the scores of the pieces they play.

Performance archive linked with scores: Users should be able to listen to all of their performances of a certain place in a score upon indicating the place in the score and selecting the recording on a timeline.

Visual performance comparison: Users should be able to visually compare differences in performance velocity among multiple recordings of the same piece. (Cf. Appendix C 3.1)

Visual performance comparison, playback integration: In cases where the users may stream the involved audio performances online, users should be able to trigger their playback by clicking on the visualization that compares them with other performances. (Cf. Appendix C 3.3)

Performance annotations: Users should be able to add textual, audio or video annotations to performances.

Exercise annotations: Users should be able to add exercise annotations to scores, which consist of a regular annotation with an addition of an example of an exercise (as audio, video or MIDI) that they recommend practicing in order to master the given musical place.

Content access control: Users should be able to set the visibility of their annotations to either private, shared with certain TROMPA users, or public.

Sharing among users: Users should be able to give access to their performances, single or all of them, with other TROMPA users.

Embeddable content: Users should be able to use the system while practicing, for example, in a conservatory, where they don't necessarily have access to a desktop or laptop computer. An application available on mobile devices, either in browser or as a native application, would be very helpful.

4.2.4 Choirs

4.2.4.1 Choral conductors

Music organization management, content access control: Users should be able to create choir organizations, appoint conductors, invite choir members to access existing repertoires.

Repertoire sharing for organization members: In order to send out the repertoire to the choir members, the system must be able to provide functionality for repertoire sharing, either by email or by creating repertoire-unique public URLs to be shared with the choir members.

Use of MusicXML scores: Conductor users should be able to create choir repertoires importing scores in MusicXML or MEI format.

Performance archive linked with scores: Users should be able to listen to all of their rehearsals of a score. Conductors should be able to retrieve and listen to performances by choir members for a particular score.

4.2.4.2 Choral singers

Member access to music organization's sheet music repertoire: Singer users should be able to access a repertoire and select any of its pieces to start practicing.

Visual representation of different voices: When selecting a piece, users should be able to visualize a piano-roll representation of the different voices and to choose one of the voices to practice. Score visualization will also be available for scores that contain score visualization information or are in MEI format (or MusicXML convertible to MEI).

Multi-language singing voice synthesis: If the lyrics or a piece are in Spanish, Catalan, Latin, English or German, users should be able to get synthesized versions in the corresponding language.

Control of vocal accompaniment: Users should be able to sing along with the synthesized voices, controlling the respective volume, pitch and panning of the different voices. (Cf. Appendix C 4.6)

Performance visualization: When singing, they should be able to get a visualization feedback of their performance (intonation, timing, etc).

4.2.5 Music enthusiasts

Visual performance comparison: Users should be able to visually compare differences in performance velocity among multiple recordings of the same piece.

Filtering compared performances by performers: Users should be able to pick particular recordings out of the list of all available recordings for comparison.

Visual performance comparison, playback integration: In cases where the users may stream the involved audio performances online, users should be able to trigger their playback by clicking on the visualization that compares them with other performances. (Cf. Appendix C 3.3)

Embeddable content: Users should be able to access to the MOOC learning contents, community discussions and analysis and annotation tools from different devices (desktop and mobile).

MOOC registration process: Users should register as participants to access MOOC contents using a web platform.

Content access control: Users should be able to set the visibility of their annotations to either private, shared with certain TROMPA users, or public.

Sharing among users: Users should be able to give access to their comments and annotations, single or all of them, with other TROMPA users.

4.3 Use-case specific requirements overlap

In this section, we illustrate the overlap between the requirements among the TROMPA use cases, as emerging from their current definitions. Requirements that are most universally shared will get first priority in implementation.

Requirement	Music scholars	Perform ers	Orchestr as	Choirs	Music enthus iasts
Embeddable content	x	x			
Content access control	x	x		x	
Search by text	x				x
Search by musical content	x				x

RDF data export	x				
Public search API	x				x
Score part generation			x	x	
Annotation propagation to score parts			x	x	
Recordings as score annotations	x				x
Score edition alignment	x				x
Practice/rehearsal statistics		x		x	
Performance archive linked with scores		x		x	
Performance annotations	x	x		x	x
Exercise annotations		x		x	
Sharing among users		x		x	x
Music organization management				x	
Mailing lists for organization members				x	
Use of MusicXML scores				x	
Member access to music organization's sheet music repertoire				x	
Visual representation of different voices				x	
Multi-language singing voice synthesis				x	
Performance visualization	x	x		x	
Visual performance comparison	x	x			x
Visual performance comparison, playback integration	x	x			x
Filtering compared performances by performers	x	x			x

5. Non-Functional Requirements

5.1 General

5.1.1 Privacy

The TROMPA project must publish a privacy policy covering all aspects of managing user data by all TROMPA member organizations.

5.1.2 Data export

Where possible, TROMPA data will be interlinked with publicly available resources and published as Linked Open Data (as RDF or JSON-LD), so as to be shareable and reusable, both across TROMPA usage contexts, and in future use-cases outside the TROMPA project.

The semantic, contextualised data serialisations thus produced promote interoperability between different kinds of TROMPA applications; for instance, by making recordings and performance-analytical metadata generated in usage contexts involving music performers available to music scholars for analysis and annotation; enthusiasts for consumption and discussion; and MIR researchers for use as machine learning corpora. Further, by making our data public in this manner, we enable reuse and exchange with external applications, making project outputs reproducible, reinterpretable, repurposable, and ultimately *useful* beyond the span of the TROMPA project.

5.2 User group specific requirements

5.2.1 Music scholars

Digital scores created by TROMPA, automatically generated links between sheet music and/or audio recordings, as well as public documents created by TROMPA users, should be exported as a Linked Data Dataset, which should include provenance information on every published relation.

The Linked Data Dataset should be versioned and published regularly, with data dumps being available as archives over HTTP on a public web page hosted by TROMPA, as well as one or more public dataset hosting platforms.

5.2.2 Orchestras

The digital scores produced by TROMPA should be of such quality that orchestras accept them as an alternative to modern commercial alternatives.

The digital scores produced by TROMPA, together with their annotations, should be exchangeable between orchestras.

Digital scores should not be exclusively playable through digital means: they should also be exportable to performance-friendly paper formats.

5.2.3 Performers

The analyses of one's performance should be available shortly after the user completed the recording.

The process of adding annotations should be as seamless as possible.

Users should be able to access practice related features on mobile phones.

5.2.4 Choirs

Users should be able to access practice related features on laptops and tablet devices.

5.2.5 Music enthusiasts

The visualization component should be embeddable by any website using a JavaScript API.

There should be a mobile implementation of a suitable subset of visualization and navigation features.

Visualizations should automatically use suitable scaling, so that useful information is conveyed for most cases.

All the tools included in the MOOC should be compatible with web browsers.

Learning contents and learning activities in the MOOC should be defined by TROMPA domain experts.

The MOOC use case should only accept participants older than 14 years old.

Participants should not require prior knowledge on the topics but they should have at least basic digital skills to understand how a MOOC works.

6. Conclusion

In this document we have presented the early application scenarios and requirements that we have devised, based on the initial feedback received from target user communities, and on the technical capabilities and content availability within the consortium.

Based on this plan, the consortium will further refine the setup of the shared contributor environment (WP5), the relation between research interests under WP3 & 4 and the present requirements, and plans towards prototypical application solutions targeted at the various user communities (WP6).

Beyond internal discussion and discussion with associated partners and the advisory board, dedicated, targeted surveys specific to each user community will be designed. To keep the user communities engaged, it is important to reach out with clear and specific questions. As a means to keep them engaged and motivated over the course of the project, we plan to keep interested users updated on our progress through the regular TROMPA communication channels.

Besides reaching out to our target user communities with refined surveys inquiring about the particular features we should develop, we are going to recruit and engage with our beta testers and frequently collect their feedback on the prototypes we are developing. For example, starting from the first half of 2019, we are going to work with students at MDW for our Performers use case. At first we will be recording data on their practicing behavior and progress, adding features over time. We will closely monitor the use of the features and seek the participants' feedback on their usefulness and usability. We also plan to work with a choir in ESMUC, on of TROMPAs associated partners. This approach will allow us to validate our assumptions and the survey feedback early, avoid unnecessary development, and create usable and useful applications.

By M18, we are going to assemble a more comprehensive requirements document that will guide the work in the remainder of the project.

7. References

7.1 List of abbreviations

Abbreviation	Description
UPF	University Pompeu Fabra
TUD	Technische Universiteit Delft
GOLD	Goldsmiths College
MDW	University of Music and Performing Arts Vienna
VD	Video Dock BV
PN	Peachnote GmbH
VL	Voctro Labs, S.L.
RCO	The Royal Concertgebouw Orchestra
CDR	Stichting Centrale Discotheek Rotterdam
MCM	Music Connection Machine
IMSLP	Petrucci Music Library (IMSLP.org)
MOOC	Massive Online Open Course
MEI	Music Encoding Initiative
RDF	Resource Description Framework
API	Application Programming Interface
GDPR	General Data Protection Regulation

Appendix A: List of closed-ended questions in the MCM survey

Here we list the six closed-ended questions (with Yes, No, and Other (a free-from answer) being the available options) of the MCM survey along with the clarifications given to the participants. The questions or statements to agree or disagree with are followed by an elaboration provided to the users in the survey form.

1. **I am interested in performance practice**

   Would you like to see a map of performance traditions for various instruments and genres and where different artists fall on it? Or an overview of multiple performances, in which you can at once see the differences in interpretations and with one click hear them?

2. **I am interested in discovering similarities among musical works**

   Would you like to look at a score and be able to see interesting connections to places in this or other scores?

3. **Would you be interested in seeing quotations and influences within and between different pieces?**

   For example, a map showing all relations between pieces (similar parts and deviations) can be shown that you can navigate, zoom and filter, where notes are highlighted and superimposed.

4. **I would like to have a map of musical influences among composers and artists**

   Imagine an animated map that shows how musical influence flowed over time, from country to country, from composer to composer, from artist to artist

5. **I would like to see what the literature says about the history and the connections between composers and/or pieces**

   We will mine lots of books about music and present the knowledge about musical connections (influences, histories, etc.) in an interactive interface

6. **Would you like to be able to easily compare sheet music editions?**

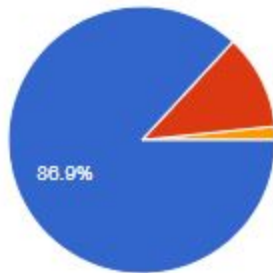   See how various editions of the same piece differ from each other

Appendix B: Summary of feedback on feature relevancy in MCM

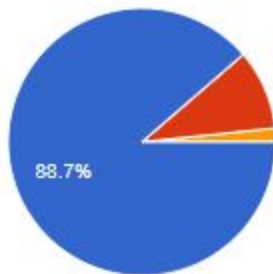**I would like to have a map of musical influences among composers and artists**
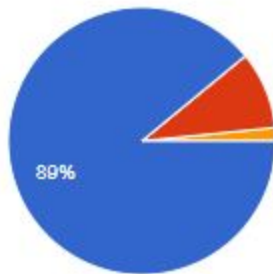
Yes	9185	89%
No	966	9.4%
Other	165	1.6%

89%

I would like to see what the literature says about the history and the connections between composers and/or pieces

Yes	8953	86.9%
No	1182	11.5%
Other	169	1.6%

86.9%

Would you like to be able to easily compare sheet music editions?

Yes	9154	88.7%
No	1001	9.7%
Other	171	1.7%

88.7%

I am interested in performance practice
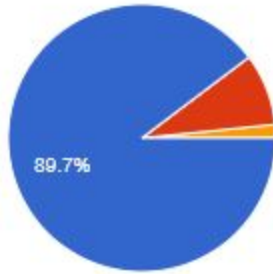
Yes	**9343**	89.7%
No	**904**	8.7%
Other	**168**	1.6%

I am interested in discovering similarities among musical works

Yes	**9498**	91.6%
No	**751**	7.2%
Other	**116**	1.1%

Would you be interested in seeing quotations and influences within and between different pieces?
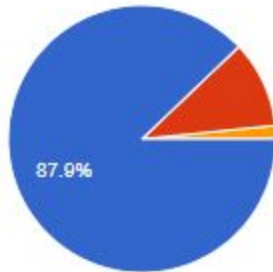
Yes	**9065**	87.9%
No	**1084**	10.5%
Other	**166**	1.6%

Appendix C: Summary of the free form MCM survey feedback

This summary outlines the most common requirements in responses to MCM survey (circa 11,600 total entries), sorted by descending popularity for each of the user groups. Quotes of the most insightful comments are provided for clarity.

1. Music scholars

1. General historical context of pieces and composer lives, linked with scores
2. Composer/Prominent musician score manuscripts with their own notes and comments. Manuscript original location information.
3. Piece similarities and influences
4. Historical musical notations, explained and "translated" to modern
5. Notations from traditions other than western
6. Possibility to search for scholarly works on the piece.
7. Identification of the source of the printed sheet music editions
8. Translations of librettos and songs to various languages.
9. The historical and geographical evolution of different instruments, linked with pieces for these instruments or arrangements for them. Technical changes that facilitated new music, eg the Boehm system for woodwinds, etc to be able to understand why music from different periods was written in that way. Data on preserved historical instruments (detailed photographs, scans, technical information).

10. Various 'dictionaries', like revealing individual melodic style of a composer

    "Brahms's sixths, for instance, or certain formulas used through centuries - leitmotiv of death (chromatic descent from T to D) in the end of Purcell's Dido and same in Bach's Crucificsus that, by the way, will help to understand the emotional content of his b-minor flute sonata (2nd element of the opening theme) and many other works by other composers where there is no verbal text, but the formula is presented"

11. Pieces comparison in different aspects.

    "Pieces that contrast in articulation (legato - non legato/staccato) or dynamics (forte - piano) or instrumentation (solo - tutti) etc. This would be interesting for teaching music theory to students and for choosing a piece for practicing purposes (e.g. student needs to learn to precisely differ legato and non legato or forte and piano etc.)"

12. Historical pedagogical works. Understanding the performance practice of past eras and the transformations in technique. They are very useful to modern performers wanting to understand how music was performed in the past.
13. Music iconography (paintings or other visual art featuring musicians). They also help understand performance practice as well as the social framing of music performance in the various eras and places.

2. Orchestras

1. Search works by instrumentation (for orchestras, also smaller ensembles, e.g. , if you want to play an orchestral piece with both a contrabassoon and a saxophone, what else can you put on the program, or can you play this piece with string quartet etc). Filter which instruments have main theme, and which can be excluded (for amateur orchestras).
2. Piece editions information and comparison
3. Make available performer/conductor data and notes about musical scores, e.g. (1) conductors' markings on scores and (2) bowings and fingerings, especially on string parts.

   "This data is now in archives or passed down from teacher to student. It could be presented as overlays on scores and parts, for study but also for production of parts to use in performance"

4. Recordings by different conductors and orchestras available for comparison with score following. Individual parts extraction (score and audio).

3. Performers

1. Comparison of different editions of a piece, as well as performance trends. Reliable information about source material, editions, publication dates

   "I recently directed a Bruckner Messe in E minor. Using all Kalmus editions, the vocal parts were different to the conductor score, which is turn was different to the orchestral parts - sometimes notes, rhythms, dynamics, added/deleted bars, inconsistent rehearsal Letter/Bar number markings etc. After much trawling the internet, it was clear that Bruckner made at least 4 separate versions across about 20 years, but determining which was which (and indeed sourcing definitive score/parts) was never entirely solved)"

2. Advanced piece search (by metadata, by melody, by musical characteristics). Ranking pieces by difficulty (with a guide to which actual aspects of the piece make it difficult: tempo, string crossings, rapid shifts, tons of accidentals, etc).

"I would be very interested in selecting a musical genre, for example a type of dance (let's say a "gavotte" of a "mazurka"), including both "titled gavottes" and especially "gavotte-like" pieces (in the sense Meredith Little and and Natalie Jenne have proposed), and then zooming my research, filtering by rhythmic pattern (e.g., gavottes including triplets in the lower metric level), by time signature, metric structure, prosodic structure of texts associated, type and main content and affect of text associated, historical period, geographical area, single composer, secondary literature associated... It would already be very interesting to start with "titled gavottes" only, but it would wonderful to include the richness of "gavotte-like" pieces, and to try to gradually define this concept more closely."

3. Watching different performances of a piece simultaneously with reading the score. Score following for performance, with annotations and tips (e.g. fingering edition suggestion by professionals, in case of string instruments)
4. Arrangements / reductions, i.e. full orchestral score for an opera, oratorio, requiem, then smaller instrumental reduction, arrangements for a single pianist and for more than one pianist. Arrangements comparison.
5. Definitions / translations (subjective) of musical terminology

    "...Including specific phrases such as Mahler's "alles ubertonend," which in the lower brass means something to the effect of 'cover/overpower every other voice in the ensemble,' whereas another instrument, composition or indeed composer may indicate a different meaning."

6. A tool that aggregates alternative fingerings and solutions for complicated instrumental parts and allows you to select between them to produce a personalised practise score for a given piece. It might be nice also to do something similar for tempo, dynamics, phrasing and ornaments.
7. Score copyright issues
8. Information alternative musical notation systems. Differences and similarities in Western and other types of music such as Indian, Chinese, African, etc. And a way to write these other musical traditions into our Western notation.

4. Choirs

1. Provide chord symbols and/or harmonic analysis of the scores to facilitate memorization / rearranging.
2. Difficulty levels for piece arrangements, for separate voices
3. Music search by note sequence
4. Piece search based on available voices (SATB, etc), info for which voices the pieces were originally written
5. Explore multiple settings of the same text (i.e. the various choral settings of "Dixit Dominus"), search by piece text.
6. Transcription/rearranging of pieces to different voices
7. Piece audio generation for non music-readers (one voice and multiple voices)

8. Canons (same idea as generated voice pieces for practice)
9. Translations