TROMPA

TROMPA: Towards Richer Online Music Public-domain Archives

Deliverable 6.4

Working Prototype for Orchestras v1

Grant Agreement nr	770376
Project runtime	May 2018 - April 2021
Document Reference	TR-D6.4-Working Prototype for Orchestras v1
Work Package	WP6 - End User Pilots
Deliverable Type	Demonstrator
Dissemination Level	PU- Public
Document due date	30 April 2020
Date of submission	29 May 2020
Leader	VD
Contact Person	david@videodock.com
Authors	David Linssen, Bauke Freiburg (VD) Marcel van Tilburg (RCO)
Reviewers	Cynthia Liem (TUD)

Executive Summary

This deliverable is the 1st version of the demonstrator deliverable for the Orchestras pilot submitted on M24 of the project. This deliverable is submitted concurrently with the other WP6 pilot deliverables that are the main deliverables related to Milestone 3 - Working integrated prototypes ready v1.0. Although these deliverables are demonstrators rather than detailed reports, the main purpose of this document is to present the functionalities of the pilot and a link to its actual implementation.

The deliverable contains three main sections. Section 2.1 presents the background of the Prototype and how the outcomes of the workshops with RCO members led to a pivot towards collective score generation instead of score annotations. The remainder of this section presents the main functionalities of the pilot by providing screenshots, a URL where we can access the pilot software, a demo video of the pilot along with instructions on how to use the pilot. We present the user interface, its different options and the available functionalities.

The first version of the prototype includes three levels of functionality: a home page for introduction to TROMPA campaigns and, in a future version, an overview of all campaigns in progress. A campaign landing page, where potential contributors can view campaign objectives, its progress, share the campaign with others and have the ability to join a campaign to contribute to the campaign. The third level is the task level, where contributors can execute tasks, depending on the objective of the campaign. Currently there are two campaigns running, contributing to the same objective. One campaign is geared towards the transcription of score snippets into MEI, the other campaign is aimed to validate those transcribed snippets. Contributors will have to consent to the publishing of their contribution under CC license and can choose to be updated on the campaign's progress via email.

A detailed description of the different incentivisation strategies are described in detail in section 2.4.2, and potential additional features that might be included in the prototype are briefly described in section 2.6.

Section 3 describes the different recruitment strategies to evaluate the prototype and to engage participation, the activities done so far for user evaluation of the pilot and the plan for the next months regarding actions to evaluate the prototype, engage participants and increase the user base. The COVID-19 crisis had an impact on all the face-to-face recruitment strategies. Thus, virtual strategies will be implemented to evaluate the prototype, while online strategies have been delayed.

Section 4 mentions the integration of WP3 technologies to the pilot, the link with WP4 human-generated data gathering strategies, and the use of the Contributor Environment and WP5 components. We observe that these integrations are in line with TROMPA Deliverable 2.2 - Complete Requirements submitted on M18. Section 5 presents the conclusions of the deliverable and future work in the context of the upcoming WP6 work.

Version Log		
#	Date	Description
v0.7	09 May 2020	Version for internal commenting
v0.8	20 May 2020	Version submitted for internal review
v1.0	29 May 2020	Final version

Table of Contents

1. Introduction

This deliverable is the 1st version of the demonstrator deliverable for the orchestras pilot submitted on M24 of the project. This deliverable is submitted concurrently with the other WP6 pilot deliverables (D6.3, D6.4, D6.5 and D6.6) that are the main deliverables related to Milestone 3 - Working integrated prototypes ready v1.0. Although these deliverables are demonstrators rather than detailed reports, the main purpose of this document is to present the functionalities of the pilot and a link to its actual implementation.

The structure of the deliverables is shared amongst all deliverables D6.3 - D6.7 and contains three main sections. Section 2 presents the main functionalities of the pilot by providing screenshots, a URL where we can access the pilot software, a demo video of the pilot along with instructions on how to use the pilot. Section 3 is dedicated to the user evaluation of the pilot. It contains information about the intended user-evaluation and the impact the COVID-19 crisis on that process. Section 4 is related to the integration into the pilot of technologies from WP3, components and methodologies from WP4, and the integration of the pilot to WP5 components.

2. Main functionalities of the prototype

2.1. Background

The main interest in the orchestras use case, as established in D6.1 'Final Mock-ups Testing', is to have TROMPA technology enable the establishment of a repository of public domain scores and orchestral parts in true digital form, including an application to access these.

Today, most orchestras have to obtain good quality, performable scores from commercial publishers, even though the musical works themselves are often in the public domain. This gives publishers power over orchestras, since publishers can exercise intellectual property rights on score editions, leading to limitations in distributing the performances further via audiovisual recordings: hiring performance materials (scores and parts) is expensive and each use other than live performance is charged extra. Secondary, current practice of publishers is based around paper distribution of scores, which creates administrative overhead for orchestra librarians and makes annotation of scores for performance cumbersome and preservation of annotations for musicology challenging.

Based on these challenges, VD designed a mock-up application for collaborative score annotations within the orchestra. The mock-up focused mainly on the score annotation interaction, based on wishes gathered from interviews with RCO members. The mock-up application then was tested during a workshop with RCO members. This resulted in three key outcomes:

- First, a set of recommendations to improve annotation interaction.
- Second, the insight that RCO members weren't likely to use such applications themselves, simply because the day-to-day practice of a top-tier orchestra wouldn't allow for it. Being a musician at RCO is a very intense job and adding a new application to the musicians' workflow would disrupt it.
- Third: all participants agreed on the practicality of a shared annotation tool for rehearsal and research. They highlighted the usefulness of such a tool for both amateur and professional musicians on the condition that there would be **a significant amount of high-quality, performable scores available** to make this digital workflow sustainable.

Currently this is not the case: most scores are available as either scanned representations of their paper originals or exist in MusicXML, a digital representation of the engraved score.Most Public Domain scores are now available as PDF (via, for instance IMSLP). Creating a digital score annotation workflow for Orchestras requires scores to be available as a digital object that could be opened and edited in score editing applications like Dorico or Sibelius, or in new applications with similar capabilities. Current MusicXML, a common interchange format, is unsuitable, because MusicXML is a representation of the music engraving, not of the musical work itself. Within TROMPA, we have therefore chosen to adopt MEI as a suitable file format. With MEI-based scores,, a number of opportunities for orchestras present themselves, such as new ways of annotating them and keeping track of these markings.

Switching to digital (annotation) workflows depends on sufficient high quality scores being available for orchestras. And without an installed base of digital score consumers, there is no incentive for publishers to deliver their library of scores in this new format.

This is a typical chicken-and-egg-type problem: the lack of high-quality, performable digital scores, because of paper-based workflows and software, blocks the adoption of digital workflows in orchestras. And without demand from performers, there is little incentive for current market players to invent and evolve their music libraries and software towards digital workflows.

2.1.1 Objective of this prototype

With the above in mind, we developed a prototype scenario where an amateur orchestra, in need of a performable score, collaborates in the conversion of a PDF-based score from the IMSLP library into a high quality MEI rendering of the same score. Such a prototype would allow us to gather insights on distributed collaboration within orchestras, add more MEI based content to the TROMPA CE library and provide an opportunity to road test the hybrid OMR functionality that TUD is developing under WP4.

Figure 2.1. TROMPA orchestras conceptual user journey.

2.2. Access information

- The **Pilot prototype** is accessible through the TROMPA subdomain[1]. The pilot prototype is currently available in English.
- A **Demo video** showing the main functionalities of the pilot can be found online[2]
- New campaigns can be set up and initiated via the backend of the software. Please contact VD for more information.

[1] https://campaigns.trompamusic.eu
[2] https://drive.google.com/file/d/1pSU8a887p-JQ2ebdIFbs4zxU18lx3AVo/

2.3. Requirements

The pilot needs the following requirements to be used:

1. A computer, tablet or smartphone.
2. A recent web browser

2.4. Pilot Functionalities

The main functionalities of the application are:

For the general public:

- Create awareness of TROMPA and specific campaign objectives
  - View TROMPA objectives
  - View running campaigns (future version)
  - Start your own campaign.
  - Contribute to running campaign.
- Provide a global overview of the campaign process

For campaign initiators:

- Setup a campaign for the MEI conversion of a PDF score
  - Add an IMSLP URL to retrieve score from IMSLP (via backend, not visible)
  - Add a campaign name, campaign objective and campaign deadline date (via backend, not visible)
  - encourage participants to collaborate on the campaign via progress emails

For Campaign Collaborators:

- Participate in Score Digitization campaigns:
  - Select a campaign to participate in (available in future version)
  - Provide a username (for named contribution purposes)
  - View tasks -> two task types currently available
  - Execute tasks -> writing MEI and MEI vs scan comparison
  - View the progress of the campaign
- Provide consent:
  - Agree with CC licensing of the created MEI file
- Enlarge the amount of contributors to the campaign:
  - Share the campaign landing page via social media share buttons or via a share link
- Configure settings:
  - Forget stored username

2.4.1. Collaborate on a campaign

The pilot core functionality is the collaboration on PDF to MEI conversion campaigns. From the Home page clicking on Join Campaign in the top menu, the user is redirected to the campaign main page. Currently there is only a single visible campaign running, where collaborators can verify snippets of MEI generated scores with their PDF originals. The details for participating in a campaign are described in Section 2.4.

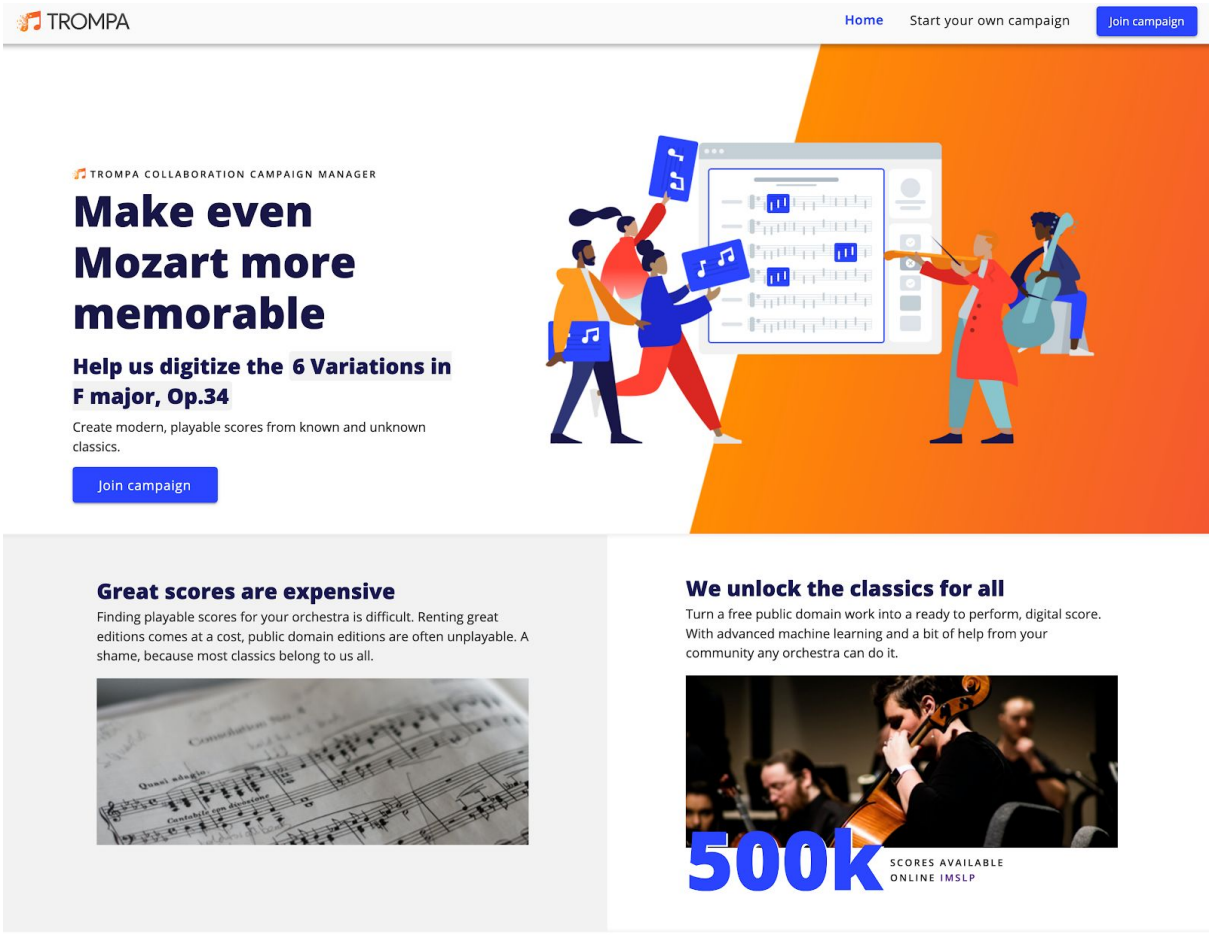

Figure 2.2. *Main* page. Users can contribute to a campaign via the Join campaign button

2.4.2. The Campaign page

For each campaign, users can visit a specific campaign landing page. On this page, they can view the campaign objectives, its progress towards completion and some general information about the work in progress. There are three main tasks on this page:

- *Joining a campaign* This will bring up the tasks pages and will present tasks to the user for completion.
- *Subscribe for updates* This will prompt the user to enter name and email address and put the user on the subscription list for a specific campaign in the email newsletter tool. From this tool we can send automated updates on the campaigns progress
- *Share this campaign* This allows the user to share this campaign with other people via buttons for Facebook, Twitter and email or by copying the link and sending it via any other platform.

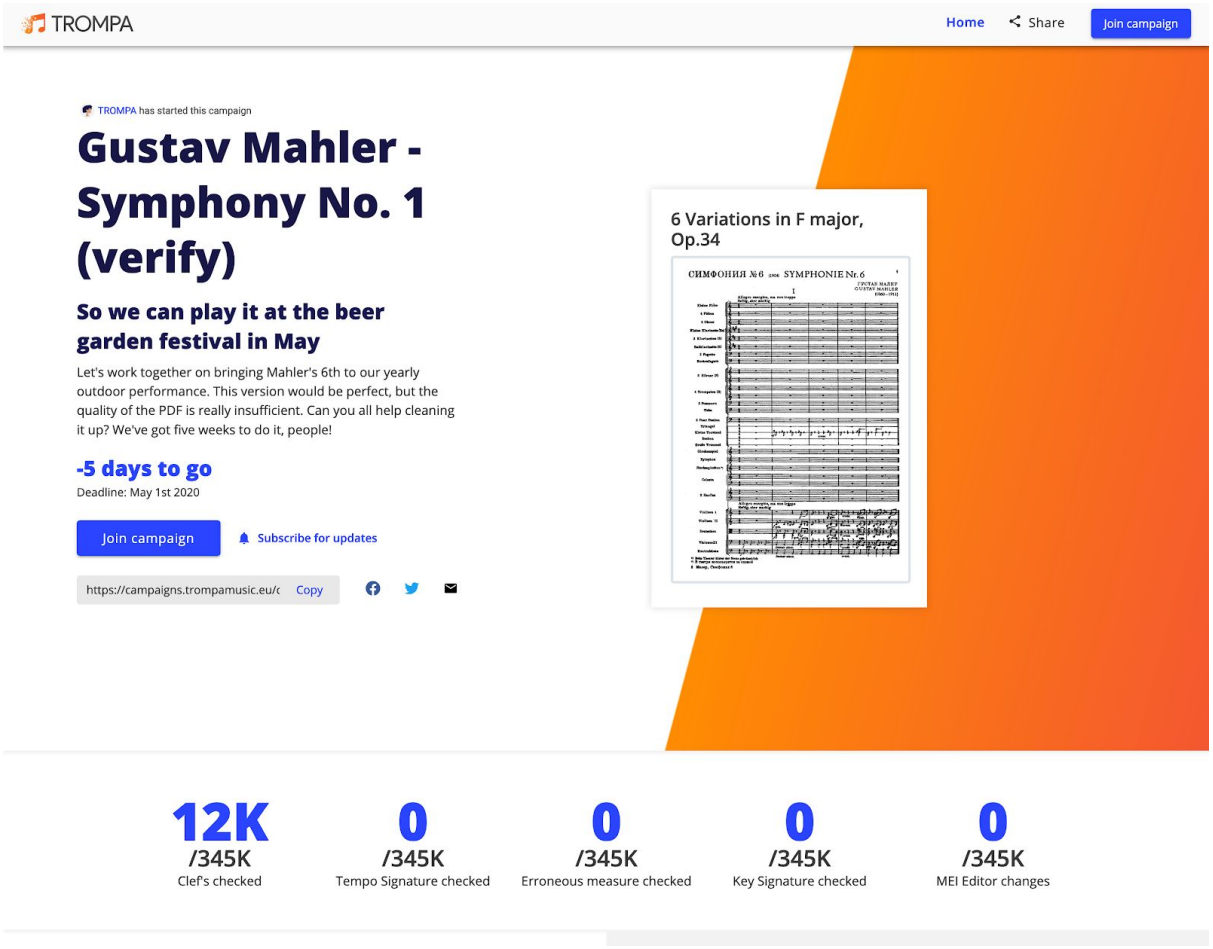

Figure 2.3. Campaign landing page. Users can view campaign objectives, subscribe for updates and invite others to the campaign via social media. The *Join campaign* button starts the distribution of tasks

2.4.3. Attribution for CC licensing

Within TROMPA, all created MEI files will be published under a Creative Commons licence (CC BY 4.0). The licence requires contributors of a work to be attributed. That's why we provide the ability to contribute under a specific nickname. Before any contribution is made to a work, users will be asked for consent to the CC licensing and will be given the opportunity to set a nickname. If users don't choose a name, their work will be published under a randomized nickname.

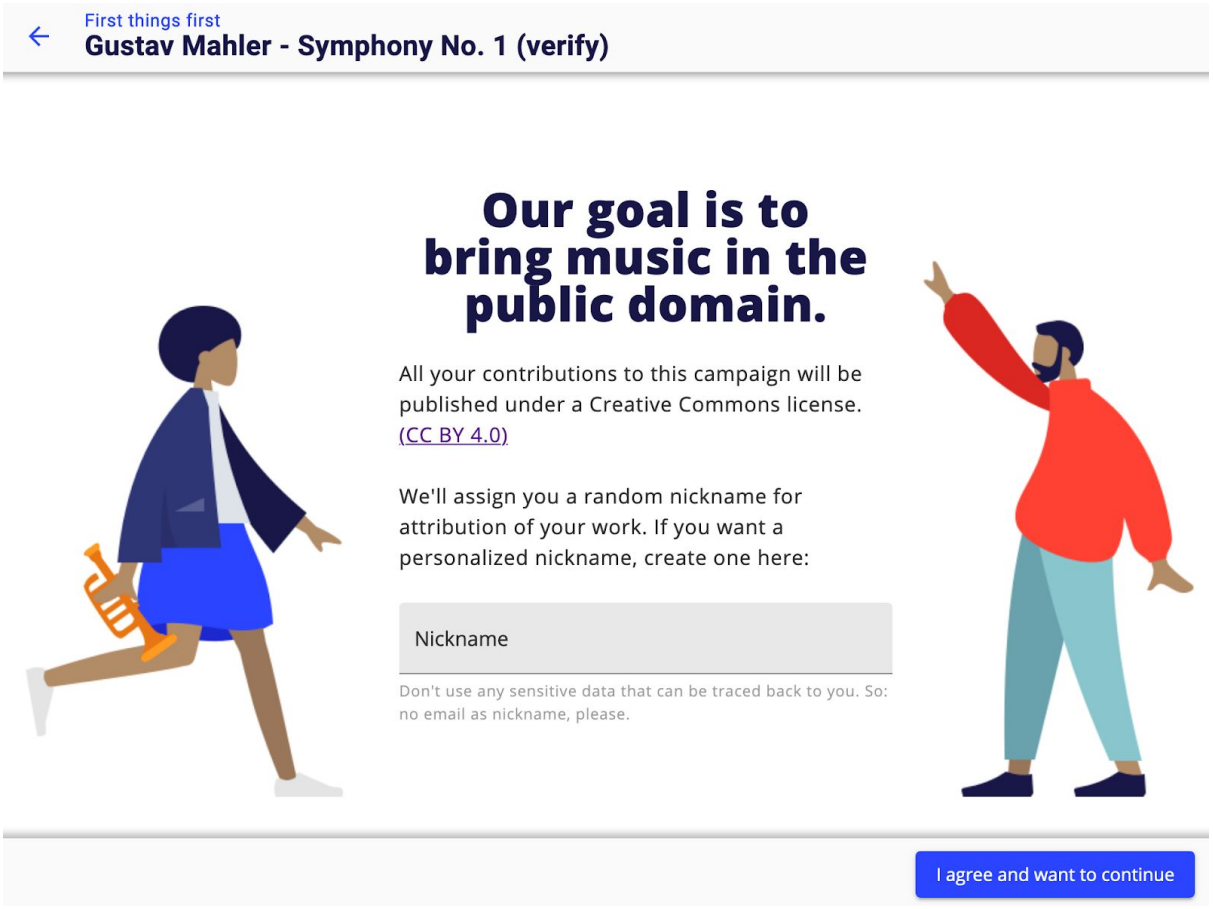

Figure 2.4. Creative Commons Consent Before contributing, users will have to agree to the Creative Commons licensing, either under a chosen nickname or a randomly assigned one.

2.4.4. The tasks page

After joining a campaign, and giving consent, the user will be served a series of tasks to complete. The nature and content of these tasks are generated by the Crowd Task Manager, a crowd-assisted Optical Music Detection (OMR) system, developed by TU Delft. In its current form, it generates transcription tasks for each measure of a given music score (see Figure 2.5), alongside validation tasks, to evaluate the output of each transcription.

In the future, it is planned to accommodate further specialised tasks, such as detecting specific music artefacts in a music score and Task Assignment capabilities, where users are receiving tasks based on capabilities.

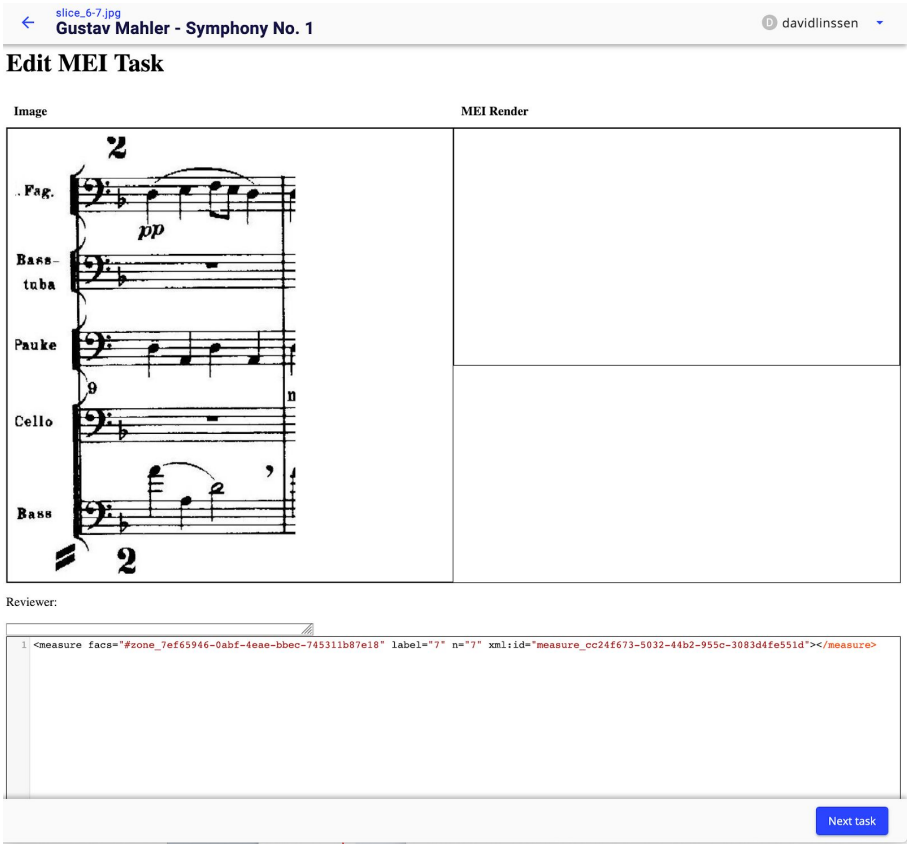

Figure 2.5. An example of an MEI transcription task within a campaign.

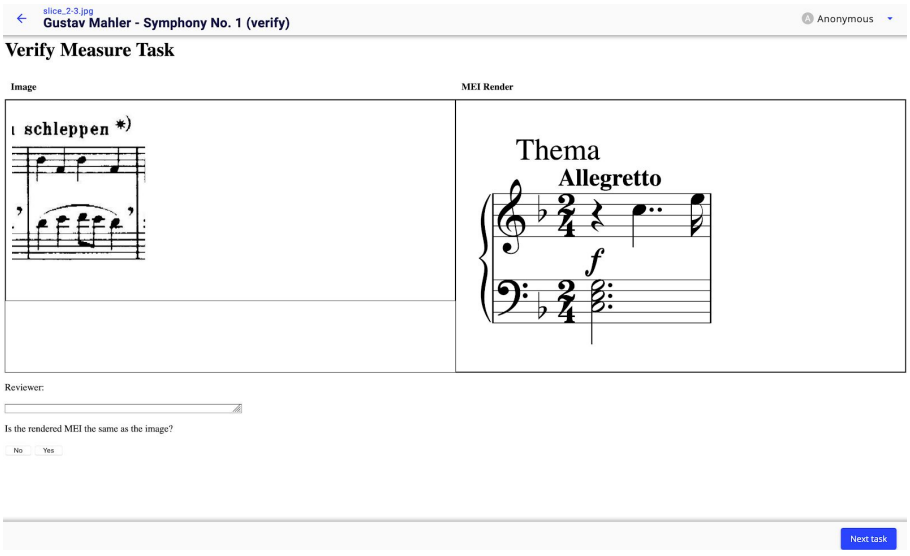

Figure 2.6. An example of an MEI verification task within a campaign.

2.5 Using the Orchestras prototype

2.5.1 Use case scenario

- The Delft students orchestra Krashna Musika wants to perform Gustav Mahler's 1st Symphony. A public domain version is available on IMSLP, however the quality of the score is insufficient for performance. The orchestra decides to use the TROMPA prototype to improve the score
- The orchestra's leader delivers the IMSLP URL to the Crowd Task Manager. This system will retrieve the score from IMSLP and generate crowdsourcing tasks per segment of the music score. For the first prototype campaign there are two task types: transcription tasks and verification tasks. The transcription tasks require participants to understand both music notation and writing XML. The verification tasks can be performed by any member of the orchestra. These tasks will be made visible in the public campaign.
- The transcription tasks can only be performed by the orchestra members that 'speak' MEI and music notation. They need to perform their tasks first, before the other members can perform verification tasks. To separate the two groups we have two campaigns running. The first campaign delivers transcription tasks to the MEI savvy group[3].
- The validation tasks are accessible to all Orchestra members via this campaign[4].
- After completion of all tasks, the resulting MEI file will be made available via GitHub and will be made discoverable through the CE -API. The orchestra can then use this digitized version for their performance.

2.5.2 How it works

The Crowd Task Manager is a system TUD built to accommodate a crowd-assisted OMR pipeline. As such it processes the input data (image processing and segmentation), it generates crowdsourcing tasks for the non-automated parts and finally aggregates results to build an MEI version of the original music score. The crowdsourcing tasks potentially can be distributed to any crowdsourcing platform of choice.

- Alongside the TUD prototype, Videodock developed such a platform to accommodate users with interest in music, namely Campaign Manager. It is a dedicated platform to facilitate crowdsourcing campaigns related to TROMPA. Through joined efforts we achieved an integration scheme between the Contributor Environment, Crowd Task Manager and Campaign Manager, as follows:
- a potential user who wants to start a crowdsourcing OMR campaign uploads a PDF file of a music score through the Campaign Manager to the Contributor Environment
- the file is fetched by Crowd Task Manager and the system generates crowdsourcing tasks per segment of the music score
- the crowdsourcing tasks are distributed to the Campaign Manager which hosts the microtasks and users can enter the platform to perform them
- the results of the crowdsourcing tasks are directly stored to Crowd Task Manager and the status of each microtask is communicated with the Contributor Environment

[3] https://trompa-campaign-manager.netlify.com/campaign/ba4314b7-2f47-4d6f-81d1-8bafb8b8f22c
[4] https://trompa-campaign-manager.netlify.app/campaign/45568fb3-e55c-48d1-9103-fc6cb33d8b04

- ❖ The Crowd Task Manager communicates with the Contributor Environment in the following fashion:
- ❖ The CE uses ControlActions (https://schema.org/ControlAction) in their GraphQL API to model campaigns and tasks in the CE. The Crowd Task Manager gets alerted once there is a new ControlAction that relates to a campaign.
- ❖ Once the pipeline has retrieved and processed the music score for crowdsourcing tasks, ControlActions are created on the CE using the RequestControlAction mutation. The resulting ControlAction is related to the campaign in the CE and contains a URL to the crowdsourcing task created by the Crowd Task Manager.
- ❖ When the Crowd Tasks Manager receives the first result from a task, the status of the related ControlAction is set to ActiveActionStatus.
- ❖ Once the Crowd Tasks Manager gets results from a task that forms a consensus, the ControlAction status is updated to CompletedActionStatus.
- ❖ For a more in-debt description of how crowdsourcing tasks are generated and the communication schemes between the above platforms, please read the 2nd version of the **Deliverable 4.4-Hybrid Annotation Workflows**[5].

2.5.3. Incentivisation strategy

In order to incentivize the contributors, we've included mechanisms as suggested in **Deliverable 4.3 - Crowd Incentivisation Mechanisms**[6].

- On the landing pages we stress the value of the campaigns.
- We make sharing easy in order to let participants use their social networks to bring in more contributors.
- We present dynamic information regarding the progress of the crowdsourcing tasks
- We keep contributors engaged through messages and announcements after completing tasks and through the campaign update emails.

Figure 2.7. In-app 'Thank you' messages keep Contributors engaged

[5] This deliverable is confidential to the consortium only
[6] This deliverable is confidential to the consortium only

2.6 Connection to requirements

The functional requirements for the pilot were described in D2.2-'Complete Requirements'. However: the feedback of the RCO participants, as expressed in D6.1 'Final Mock-ups Testing' , underlining their extremely high output standards and declaring the use of such a system as too much overhead for musicians, and the lack of performable scores in MEI format have led to a pivot in the past year. As mentioned in D6.2 'Planning for the execution of pilots in real life settings', contacts have been established with amateur student symphony orchestras in The Netherlands. This allows for the establishment of a larger engaged orchestra community, that is more dependent on the availability of playable digital scores, and that is digitally minded.

The insights from RCO have led us to investigate the needs of other (amateur) orchestras. The realities of an amateur orchestra are totally different from the needs of members of a top-10 global orchestra. Both groups indicated that - given the choice - they would prefer to annotate their parts the old-fashioned way, with pencil, which is in their view faster and more practical.

At the same time: both groups underlined their need for good quality, performable scores in the public domain.

Currently, the only feasible way to make more scores available in MEI format is by encoding them by hand. This is a very laborious process which can only be performed by a skilled musicologist. To speed up the process we need to look at machine-human hybrid systems. With this pilot prototype we hope to gain more insights on human-hybrid collaborations in practice and build a working application framework which we can use for the testing of more powerful recognition algorithms. As part of efforts under WP3, T3.4, TUD and PN are working on automated vision-based strategies to pre-populate and help prioritizing score fragments to be corrected.

To make the Campaign Manager as reusable as possible, it is built as a separate component and can handle task types of any kind.

2.7 Future outlook

We aim to connect the prototype to the Multimodal Component, so it will be able to point users to multiple campaigns. We also aim to use the MMC to let Campaign Leaders select scores from IMSLP to start the transcription process . We aim to incorporate other MEI editing components in tasks. Ideally we come to a situation where MEI writing skills won't be necessary, as graphical interface improvements will help reduce the cognitive load of users.

3. User evaluations

3.1. Initial target audience & recruitment strategies

The initial planning of target audience and recruitment strategies had been as following:

- Mid-April 2020: conduct first user tests with Krashna Musika.

3.2. Impact of COVID-19 crisis

Due to COVID-19 crisis, the test audience, Krashna Musika, can not rehearse. Some members have returned to their home countries, for others continuing their studies during the quarantine period has proven challenging. All in all, the willingness to participate is still there, but setting up the test has been proven challenging. We aim to conduct the user testing before the end of the academic year on July 3rd, but given uncertainty in student availabilities, we also are researching collaboration opportunities with other amateur ensembles, and may need to push this test into July.

3.3. Adjusted target audience & recruitment strategies

Our target audience has not changed, but the recruitment strategies have been adjusted, as follows (as long as face to face activities are restricted):

- We will conduct a preliminary test with members of the TROMPA consortium to end-to-end test the prototype in May
- We will set-up an online test with amateur orchestra players (preferably Krashna Musika, but with back-up from other ensembles) no later than July.

3.4. User evaluation outcomes

3.4.1 Online evaluation study protocol

- Not yet available

4. Current and future integration with other TROMPA WPs

This section summarizes the integration of the prototype with the components and technologies developed in other WPs, including WP3 (automatic description), WP4 (crowd-sourcing and human computation technologies) and WP5 (contributor environment and core components).

4.1. Relation with TROMPA WP3

WP3 tasks	Currently integrated	To be integrated in next version
Music description	n/a	n/a
Audio processing	n/a	n/a
Visual analysis of scores	n/a	X
Alignment of musical resources	n/a	n/a
Multimodal cross-linking	n/a	n/a

4.2. Relation with TROMPA WP4

WP4 tasks	Currently integrated	To be integrated in next version
Crowd-powered improvement	X	X
Annotators	n/a	X
Incentivisation of TROMPA crowds	X	X
Campaign design	X	X

4.2.1 Crowd-powered improvement

The entire objective of this Prototype is to gain a better understanding of Crowd-powered improvements. That's why the prototype was built with the ability to address a large group of diverse people in mind. Once we have sufficient experience with the prototype in a controlled environment, we aim to open up the platform to a larger, general audience.

4.2.2 Annotators

This prototype is built making use of and conforming to **Deliverable 4.2 - Annotator Properties and Metrics**[7]. In this first version of the prototype, we target users whose skills are known (e.g. orchestra members) thus their Competence Model is assumed. But it's not yet possible to classify an unknown crowd according to the Competence Model, as user-related information is not being registered yet on the Contributor Environment.

4.2.3 Incentivisation of TROMPA crowds

We have implemented many recommendations made in the 2nd version of **Deliverable 4.3 - Crowd Incentivisation Strategies**, as mentioned under 2.4.2.

4.2.4 Campaign design

This prototype will allow us to experiment with multiple strategies for attracting, engaging and rewarding contributors or crowd workers.

4.3. Relation with TROMPA WP5

WP5 components	Currently integrated	To be integrated in next version
Score edition component	n/a	X
Processing library	n/a	n/a
Multimodal integration	X	X
Performance assessment	n/a	n/a
Annotation tools	n/a	X

4.3.1 Multimodal integration

The prototype relies on the ControlActions mechanism of the TROMPA CE for the communication between the Crowd Task Manager of the hybrid annotation pipeline described in **Deliverable 4.4 - Hybrid Annotation Workflows** and the user-facing CampaignManager. A detailed description of this mechanism can be found in paragraph 5.1 of the D4.4 document.

The Multimodal Component will be integrated in a future version of the prototype, to allow users to navigate between multiple running campaigns.

4.3.2 Score edition component and Annotation tools

We're investigating the feasibility of integrating the Annotation Tools or the Score Edition Components in a future version. These would allow crowd workers/collaborators to perform more specific tasks than in the current version.

[7] This deliverable is confidential to the consortium only

5. Conclusion

The current version of the Orchestra Prototype is geared towards the collective creation or improvement of public-domain musical scores. It combines a collaborative campaign mechanism with the OMR capabilities of the Hybrid Annotation Workflows. The prototype will allow us to gather some real-world test data on actual collaboration processes within amateur orchestras, which we can use to refine the hybrid annotation workflow.

In its current state, its functionality to create performable MEI is very limited. The OMR engine is able to split entire scores in microtasks that can be presented to the crowd workers, to validate their input by presenting it to other crowd workers and recombine the overall results into an MEI file. However, the conversion of these image snippets to MEI is still a manual process. In order to get MEI output as scale, the actual recognition process needs to be performed by computers, rather than humans.