Project Acronym:	**Europeana Sounds**
Grant Agreement no:	**620591**
Project Title:	**Europeana Sounds**

D2.4 Crowdsourcing infrastructure V1 Assessment and Recommendations

Revision: Final

Date: 30/06/15

Authors: Maarten Brinkerink (NISV), Antoine Isaac & Hugo Maguinhas (EF), Harry Van Biessum (NISV), Vassilis Tzouvaras (NTUA)

Abstract: We provide an update on the current state of the development of the crowdsourcing infrastructure, as well as some initial evaluation outcomes resulting from user research conducted for some of the end-user facing components (crowdsourcing applications) of the system architecture.

Dissemination level	
Public	X
Confidential, only for the members of the Consortium and Commission Services	

Coordinated by the British Library, the Europeana Sounds project is co-funded by the European Union, through the ICT Policy Support Programme as part of the Competitiveness and Innovation Framework Programme (CIP)
http://ec.europa.eu/information_society/activities/ict_psp/

Revision history

Version	Status	Name, organisation	Date	Changes
0.1	ToC	Maarten Brinkerink (NISV)	03/06/15	
0.2	1st draft	Maarten Brinkerink (NISV)	09/06/15	First draft, based on MS11
0.3	Draft	Antoine Isaac & Hugo Maguinhas (EF)	15/06/15	Review from EF R&D Perspective
0.4	Draft	Harry Van Biessum (NISV)	26/06/15	Chapter 3
0.5	Draft	Technical coordination group members (NISV, EF, AIT, NTUA, NET7, SHIFT)	26/06/15	Added reflection on the executed work and lessons from the Athens workshop
0.6	Draft	Vassilis Tzouvaras (NTUA)	28/06/15	WITH modules added to chapter 1
0.7	Draft	Segiu Gordea (AIT)	29/06/15	Updated annotations API information
0.8	Final Draft	Maarten Brinkerink (NISV)	30/06/15	Final version delivered to BL
1.0	Final	Laura Miles, Richard Ranft (BL)	01/07/15	Layout, minor changes

Review and approval

Action	Name, organisation	Date
Reviewed by	Ashley Burgoyne Remy Gardien	15/06/15 30/06/15
Approved by	Coordinator and PMB	01/07/15

Distribution

No.	Date	Comment	Partner / WP
1	02/07/15	Submitted to the European Commission	BL/WP7
2	02/07/15	Posted on Europeana Pro website	BL/WP7
3	02/07/15	Distributed to project consortium	BL/WP7

Application area

This document is a formal output for the European Commission, applicable to all members of the Europeana Sounds project and beneficiaries. This document reflects only the author's views and the European Union is not liable for any use that might be made of information contained therein.

Statement of originality

This document contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.

Project summary

Europeana Sounds is Europeana's 'missing' fifth domain aggregator, joining APEX (Archives), EUscreen (television), the Europeana film Gateway (film) and TEL (libraries). It will increase the opportunities for access to and creative re-use of Europeana's audio and audio-related content and will build a sustainable best practice network of stakeholders in the content value chain to aggregate, enrich and share a critical mass of audio that meets the needs of public audiences, the creative industries (notably publishers) and researchers. The consortium of 24 partners will:

Double the number of audio items accessible through Europeana to over 1 million and improve geographical and thematic coverage by aggregating items with widespread popular appeal such as contemporary and classical music, traditional and folk music, the natural world, oral memory and languages and dialects.

Add meaningful contextual knowledge and medium-specific metadata to 2 million items in Europeana's audio and audio-related collections, developing techniques for cross-media and cross-collection linking.

Develop and validate audience specific sound channels and a distributed crowd-sourcing infrastructure for end-users that will improve Europeana's search facility, navigation and user experience. These can then be used for other communities and other media.

Engage music publishers and rights holders in efforts to make more material accessible online through Europeana by resolving domain constraints and lack of access to commercially unviable (i.e. out-of-commerce) content.

These outcomes will be achieved through a network of leading sound archives working with specialists in audiovisual technology, rights issues, and software development. The network will expand to include other data-providers and mainstream distribution platforms (Historypin, Spotify, SoundCloud) to ensure the widest possible availability of their content.

For more information, visit http://pro.europeana.eu/web/europeana-sounds and http://www.europeanasounds.eu

Copyright notice

CONTENTS

Executive summary: D2.4 Crowdsourcing infrastructure V1 Assessment and Recommendations

This document provides an update on the current state of the development of the Crowdsourcing Infrastructure, as well as the evaluation outcomes resulting from user research conducted for some of the end-user facing components (crowdsourcing applications) of the system architecture.

The first section describes the state of the Crowdsourcing Infrastructure development. To start it provides an updated system architecture overview (See: Image 1), to illustrate the advancement compared to the design that was provided in *D2.2 Functional design of semantic enrichment* (Ref 1). Then this section goes on the describe the development status of the crucial elements, the Core Infrastructure (1.1), the Annotations API (1.2), the Europeana Channels (1.3), and the WITH modules (1.4) in more detail.

The second section describes the advancement of Shift with the *Crowdsourcing through specialised platforms scenario* (as introduced in D2.2), which focusses on the challenge of music identification. It describes the underlying user research, the proposed technical solution, first progress with the technical development, and a timeline for the development and integration.

The third section describes the advancement of NET7 with the *Manual correction of semantic annotations* scenario (as described in D2.2), which focusses on enhancing existing metadata records through a crowdsourcing widget based on Pundit, and the evaluation of the current state of this work (as executed by NISV). It describes the development of a pilot version of the crowdsourcing widget, user research that has been applied to this pilot version, describes the problem areas that emerged from the evaluation (including recommendations for improvement), and finally looks ahead at the upcoming work, including the next pilot scenario.

The fourth and final section summarizes the lessons-learned and recommendation resulting from the work conducted in the context of the Crowdsourcing Infrastructure up until now.

1. System architecture update

This first section describes the state of the Crowdsourcing Infrastructure development. To start it provides the finalized version of the updated system architecture overview (See: Image 1), to illustrate the advancement compared to the design that was provided in *D2.2 Functional design of semantic enrichment* (Ref 1) and *MS11: Evaluation of first deployment of the crowdsourcing infrastructure* (Ref 2). Then this section goes on the describe the development status of the crucial elements, the Core Infrastructure (1.1), the Annotations API (1.2), the Europeana Channels (1.3), and the WITH modules (1.4) in more detail.

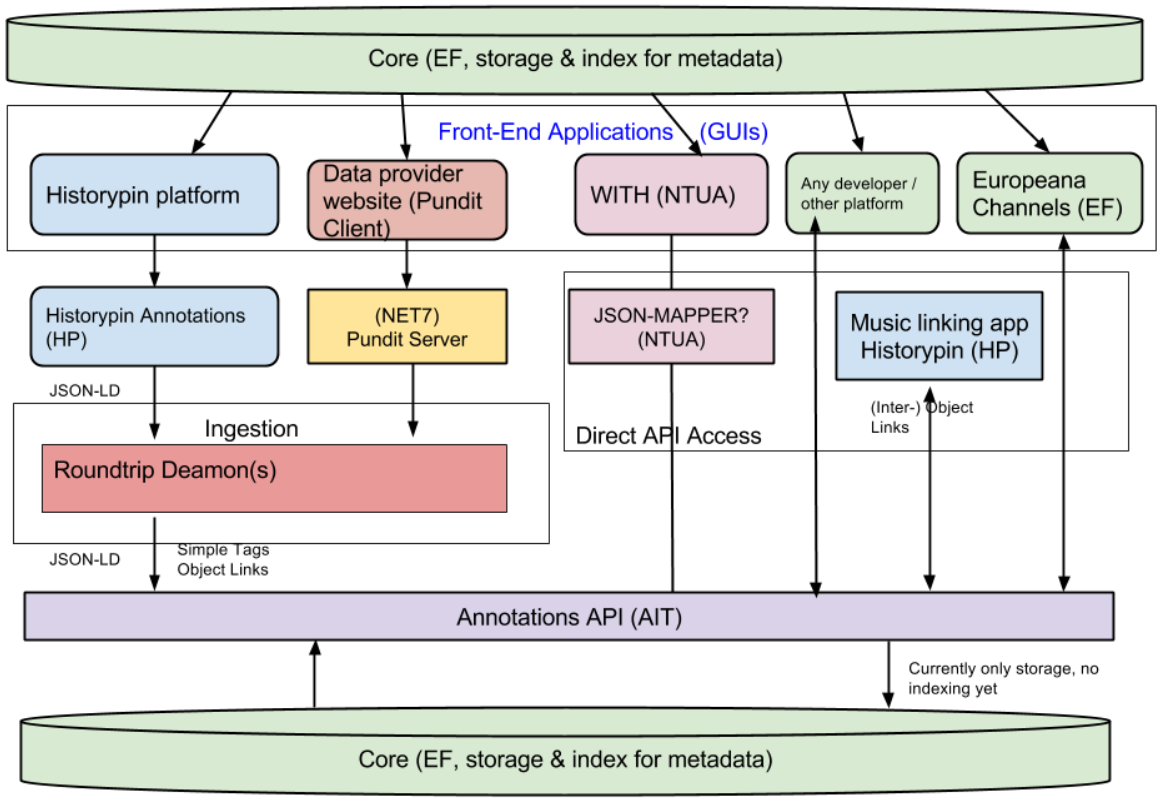

Image 1: Updated system architecture overview

1.1 Core infrastructure

As the Europeana infrastructure is the one designated to aggregate all Europeana Sounds content, and in particular all data related to enrichments gathered through crowdsourcing, the first focus for this work has been to extend its core with the ability to ingest and process these enrichments.

Europeana's institutional metadata is manually provided by Data Providers, and is then made available via monthly publications on the Europeana Portal and via the Europeana API services.

Crowdsourced data is not as manageable as institutional metadata is. It requires a different workflow. Institutional metadata can be processed on a monthly basis, while crowdsourced data can

come in at any date and time, and cannot wait for a scheduled timeframe to be processed and made available.

This has led to the following tasks for EF to start adapting its infrastructure to be able to process crowdsourced data such as annotations and enrichments:

1. Integrate the Europeana Annotations API, as developed by AIT (see section 1.2), into the Europeana infrastructure to have one single point of entry for crowdsourced enrichments/annotations.
2. Set-up database infrastructure to make crowdsourced enrichments available.
3. Adapt Europeana's Ingestion workflow in order to process and publish data outside of the specified scheduled intervals, in particular when it comes to indexing requirements.

The first activities that have taken place in the context of Europeana Sounds were focussed on tasks 1 and 2. Further work has been planned to gather and define requirements that will function as a basis for the work that should be conducted in task 3.

To align on the integration of the various components within the Core infrastructure, an official workshop was planned and took place around the Training Days programme in Athens end of June, 2015. Some of the outcomes from this workshop are included in this document.

1.1.1 Integration of the annotations API

EF is working closely with AIT to further develop its Annotations API, which will serve as a basis to process incoming crowdsourced enrichments in the form of annotations. We have currently integrated a first version of the Annotations API into the Europeana infrastructure, which serves as a pilot to process incoming annotations via selected platforms, as well as a basis to further iterate on with the ultimate goal being to make the Annotations API available for public use.

Alongside, this also brings a need to adapt Europeana's core API to be adapted to crowdsourcing. An example of this is when a Europeana record is retrieved via its API, the crowdsourced data should also be included in the response. The same will apply for search.

1.1.2 Round-trip daemon

The first version of the Annotation API will provide support for the implementation of the tagging and object linking epics (as detailed in D2.2). The definition of the API methods was made in the larger context of transferring Annotations from the Historypin[1] platform into the Annotation API and it is available in these documents: Round Tripping Daemon[2] and API Specifications for HistoryPin Pilot[3].

As part of Europeana v3, Shift has been working on the two way flow of content and metadata between Europeana and the Historypin platform. For Europeana, a strategic goal is to be able to act

[1] See: http://www.historypin.com/
[2] See: https://docs.google.com/document/d/1eP6CVFLtKqSCduhcuYnqPEKQR7_Za0cizNYE0SQ-6Og
[3] See: https://docs.google.com/document/d/15DaWKOHHxqxv4QxNi4wHxCbVzfnJu8nUX8GUK_mQ5ek

as a hub for metadata and digital objects, where the flow of content is two way and as automated as possible. Shift is the first project partner to pilot this two way content flow. The overall workflow is:

1. Historypin calls data from Europeana via its API and imports this into the Historypin platform, in order to create a digital object on Historypin for each object from Europeana
2. Historypin users add semantic enrichments (i.e. annotations) to these digital objects on Historypin.org
3. Europeana calls Historypin.org and imports the annotation data into Europeana

Currently, the Europeana Round Tripping Daemon is a prototype which can read these specific two types of annotations from Historypin and import them into Europeana. In the coming year, the Round Tripping Daemon will either be replaced by different software, or be upgraded, to accommodate also other providers and types of annotations. All of this with the goal of providing a generic and automated way for importing annotations into the Europeana platform.

1.1.3 Set-up infrastructure for crowdsourced enrichments

As well as the need to deploy the Annotations API on Europeana servers, the crowdsourced enrichments must also live somewhere. For the first iteration of the work, the crowdsourced data will be stored in a MongoDB. It is a separate MongoDB collection that lives outside of Europeana's institutional metadata. This ensures there is a clear separation between the user-contributed metadata and the original metadata record provided by the institution. The reference that binds the two together will be the Europeana Object ID.

1.1.4 Adapting the ingestion workflow

As described in the first paragraphs of this chapter, the nature of crowdsourced enrichments calls for a change in Europeana's Ingestion workflow to be able to process and publish this data outside of the regular scheduled content publication of Europeana.

The biggest challenge here are the indexing needs: ideally, when a user makes an annotation, this annotation instantly becomes searchable and available to other users. This obviously cannot take weeks, but must happen closer to real-time. This poses a challenge as Europeana's search index (Solr) is currently only updated on a monthly basis.

Together with AIT and the Technology Team at EF, discussions have been initiated to answer two questions:

1. What are the exact indexing requirements when it comes to annotations and enrichments? How long may it take for a user to find an annotation, and how should crowdsourced enrichments be integrated into the 'regular' Europeana search?
2. How can we update the Europeana index while still being able to process and update institutional content?

In June, Europeana organised a technical workshop for Annotations together with AIT, and part of it

also attempted to solve these questions. This has led to the following remarks and recommendations. On indexing requirements:

- Annotations must be searchable almost instantly after they are added, as part of the Europeana main search (thus not as a separated annotations search).
- The time it takes for an annotation to be indexed must be as short as possible and may not be longer than 15 minutes.
- An example of an annotation which has impact on the Europeana search is the scenario where:
    - A user adds the tag 'church' to a record about a church.
    - Another user searches on 'church'.
    - Europeana search then includes this record in the search results.
- All metadata on annotations which directly relates to the Europeana main search must be stored in the Europeana main index, as part of the record.
- All other metadata on annotations which should be searchable as part of annotations (e.g. retrieve all annotations by a specific user) should be part of a separated annotations index and a separated search.

On the second question to answer, on being able to update the Europeana index while still being able to process and update institutional content, progress is made in the Technology team of Europeana towards a process of continuous publication: being able to update the index at all time from different systems and tools. This however, is a big effort; hence the aim for this to be completed and integrated is no later than January 2016.

1.1.5 Other relevant development progress

Other work that has taken place with Europeana that has a relevance to the work on the Crowdsourcing Infrastructure:

- To support the crowdsourcing scenarios (as defined in D2.2), as well as the Europeana Channels project (WP4), adaptations must be made to the Europeana Search API. The focus in the first months has been to further mature the API, optimise its performance and provide a better way for testing API responses. The next iterations will focus on adapting and introducing API methods based on feedback and discussions within WP2 and WP4.
- To retrieve annotations from Europeana at very large scale, the Europeana REST-API might not always be the best fit for purpose. Europeana is therefore working on making an OAI-PMH Data Provider available for third parties to harvest Europeana data in bulk. This should be available as a beta service in the next quarter (Q3, 2015).
- Work is taking place in the Europeana v3 project to import annotations on Europeana objects that have been created on the Historypin platform. This a first implementation of the Europeana Annotations API and includes a custom prototype of a roundtripping application which maps the Historypin API output to Europeana Annotations API input.

Furthermore, work is being done on the topic of data modelling to define a generic way to model

annotations/crowdsourced enrichments when it comes to Ingestion and API services[4]. These collaborated efforts are taking place within Europeana in general, within Europeana Sounds, together with various partners.

1.1.6 Lessons-learned and recommendations

When developing a set of crowdsourcing tools of which the majority will be integrated into one coherent infrastructure, it is vital to align all efforts in every stage of the project. This refers to things such as:

- Aligning all the work on data modelling and specifications in order to make sure data can be exchanged with the least technical efforts as possible (e.g. using a standard format for transferring annotations between Pundit[5] and the Annotations API).
- Working from the start with an integration perspective and iterative development (releases) in mind: test the integration at every stage. Europeana can provide guidelines and test environments to work against.

Part of the work under the Taskforce on Enrichment and Evaluation Methods[6], organized in the context of EuropeanaTech (and with WP2 being represented by WP lead Maarten Brinkerink, as reported in D2.2), has been the creation of an annotated corpus of metadata records (also known as 'Gold Standard', Ref 3), against which the results of an automatic enrichment tool can be evaluated to determine its qualitative performance. While creating this corpus, the task force recognized that the task of annotating the corpus, and also determining the correctness of the annotations, could be a relevant application and testing scenario for the Europeana Sounds Crowdsourcing Infrastructure. Looking at the Epics described in D2.2 (Ref 1), it makes sense to enhance the *Object tagging* epic (E2) with a specialization to support semantic enrichments (following the proposal[7] from the PATH project[8] to model semantic enrichments as annotations using the Open Annotation model, which has been adopted by WP2), and *Moderation epic* (E9) for reaching the agreement and confirmation that the annotation is deemed correct.

A possible setting could be the use of the Pundit client from NET7 on-top of a Europeana Portal record view to expose the corpus' annotations, also allowing the identification of new annotations and ability to determine their value through moderation, while the Annotations API would support the persistence of the Gold Standard.

Europeana and AIT organized a technical workshop on the Annotations API (which is described below) in June 2015, with the goal of creating a stable design of the RESTful API. The following topics were discussed:

- API essentials (such as data formats, naming conventions)

[4] See: https://docs.google.com/document/d/1I5pFY3WqLn83_mWXxbXGt9_eT3-VRcfkygZ-Ag-k2Io/
[5] See: http://thepund.it/
[6] See: http://pro.europeana.eu/evaluation-and-enrichments/
[7] See: http://www.paths-project.eu/eng/content/download/5342/40580/version/6/file/edmPATHS.pdf
[8] See: http://www.paths-project.eu/

- API methods (naming, parameters, use-cases)
- Indexing and search requirements

During the meeting we drafted an updated specification for the Annotations API and agreed on the use-cases and functional requirements to be developed and implemented as part of the project. We decided to follow a roadmap that creates a new release of the Annotation API each 6 months starting with July 2015. A first draft of the outcomes of this meeting will be circulated soon among the (technical) partners in the project for feedback. The implementation of the annotation round trip with HistoryPin and Pundit will have the goals to stabilize the data model for representing the Annotations. For the specification of different types of annotation the W3G working group on Web Annotation is also consulted to discuss specific JSON-LD representation details.

1.2 Annotations API

The Annotations API is a backend application implementing functionality for the storage, retrieval and management of annotations. The development of this service was started within the scope of the Europeana Creative project, following the Open Annotation specifications[9] and supporting JSON serialization. Within the scope of the Europeana Sounds project, the service was enhanced to support the new standardized and more flexible representations of data in the linked data format (JSON-LD) specified within a draft version of the new Web Annotations[10] standard.

The first release of the JSON-LD based API is in an acceptance testing phase and will be deployed as part of the Crowdsourcing Infrastructure. Currently it supports the storage and retrieval of tagging and object linking epics (as included in D2.2). The functionality regarding the indexing, retrieval and administration of annotations is work in progress and will be released as part of the next deployment to production, later this year (Q3 or 4). Below the main new API methods will be described.

The Annotations API was also enhanced with the integration of a standardized application console based on Swagger[11]. This provides support for creating standardized documentation and testing of the REST API. The Alpha version of the Annotation API[12] was deployed within the Europeana test environment for testing and integration evaluation purposes. The next version of the Annotation API is planned to be released in January 2016 and will include support for semantic tagging, image annotations, indexing and retrieval and basic functionality for moderating the annotation lifecycle.

1.2.1 Create annotations - API method

The creation of different types of annotations uses the same API method, but it uses different validators for ensuring consistent input data. The specifications of the newly released method are available in the aforementioned *API specification for Historypin pilot* document. The examples

[9] See: http://www.openannotation.org/spec/core/
[10] See: http://www.w3.org/TR/2014/WD-annotation-model-20141211/
[11] See: http://swagger.io/
[12] See: http://annotation-web.de.a9sapp.eu/docs/index.html#!/europeanald

provided in this document were integrated as templates[13] in the Swagger documentation (see implementation notes description).

The input parameters for this method include the wskey, which is reserved for authentication and authorization purposes. The annotation numbers might be provided as parameters, or generated automatically by the system, according to the configuration policies. The annotation content itself is provided within the request body using JSON-LD serialization. Future releases are planned to add separate methods for creating annotations basing on the motivation used for their creation (e.g. tagging, linking, rating, commenting, etc.). The compact JSON-LD representation of these different types of annotations are expected to be different from each other, but also is expected to be the preferred way of creating annotations, while the verbose representations(e.g. including the whole provenance information) are expected to be used for moderation and administration purposes.

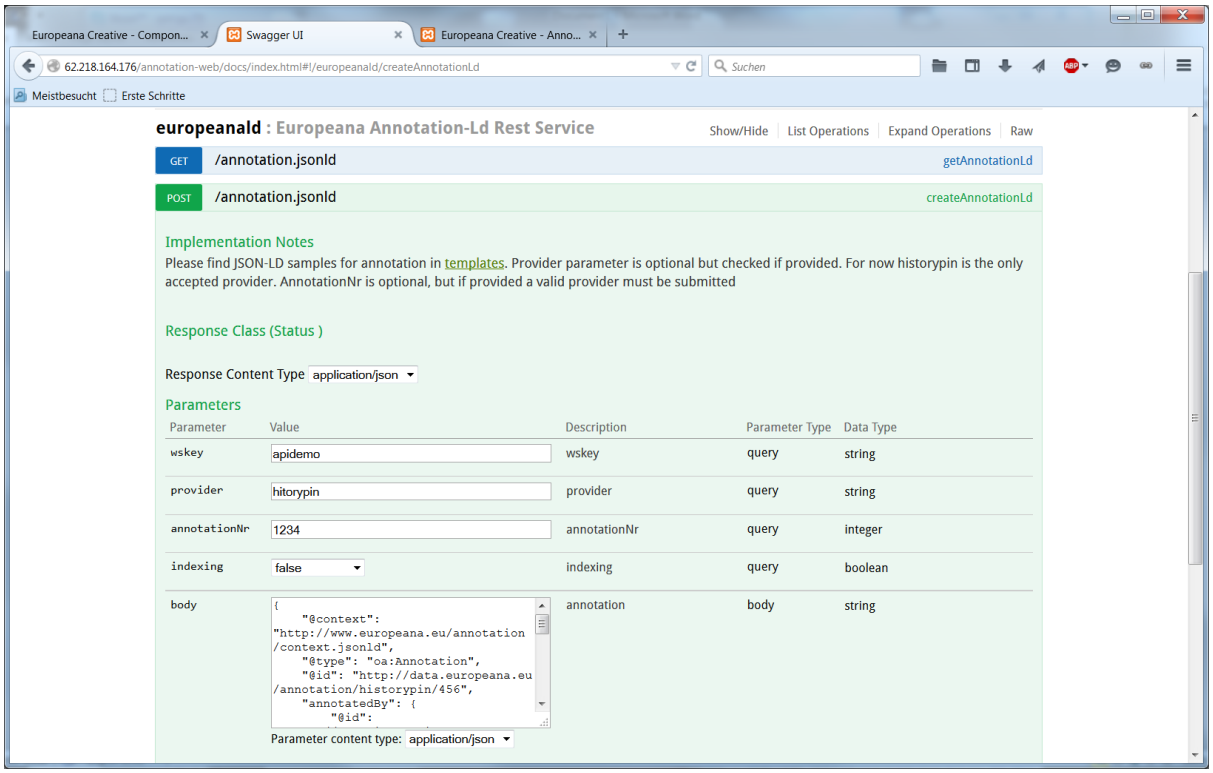

Image 2: Annotations Console - create annotation

1.2.2 Read annotations - API method

This API method provides access to the annotations stored within the database, basing on their id composed from the combination of a provider name and an annotation number. The annotations are represented similar to the input method; the only additional field that is returned by this method is the annotation uri, which is available in the @id field.

[13] See: http://62.218.164.176:8080/annotation-web/docs/templates.jsp#europeana-ld-api

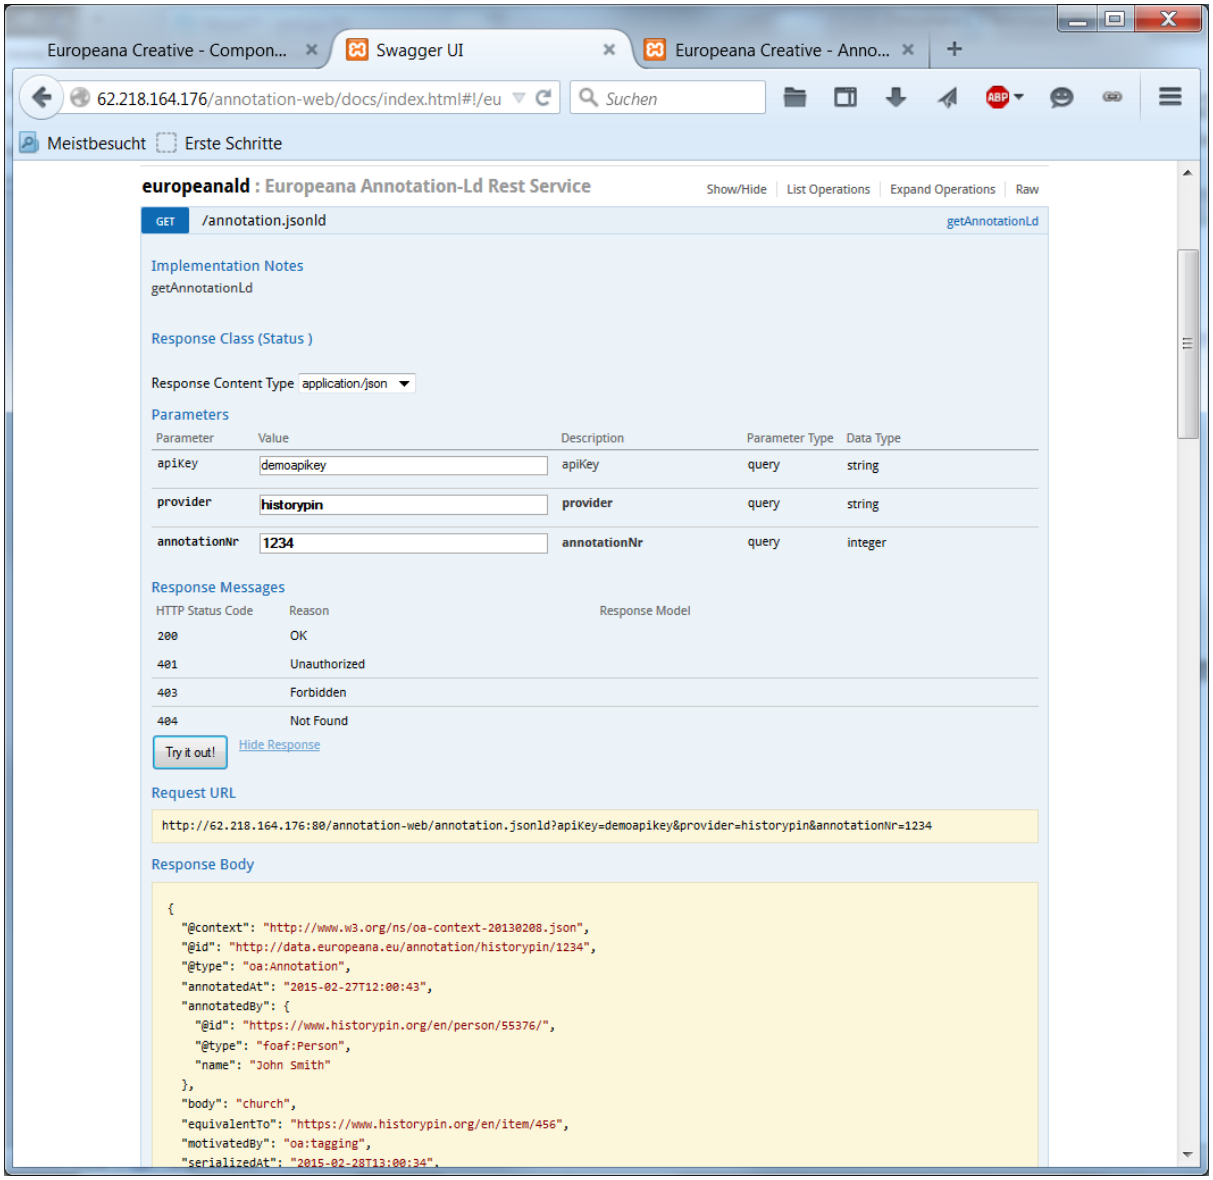

Image 3: Annotations Console - read annotation

1.2.3 Annotation search and administration APIs

The retrieval and administration API methods are in a prototype status and are planned to be released in the next versions of the API. These admin interfaces provide access to functionality that disables, deletes and/or reindexes annotation objects, but also for the deletion of the concepts used for tagging objects (i.e. “delete tags”). The search interface provides methods that perform a free text and/or metadata search for retrieving annotations and tags.

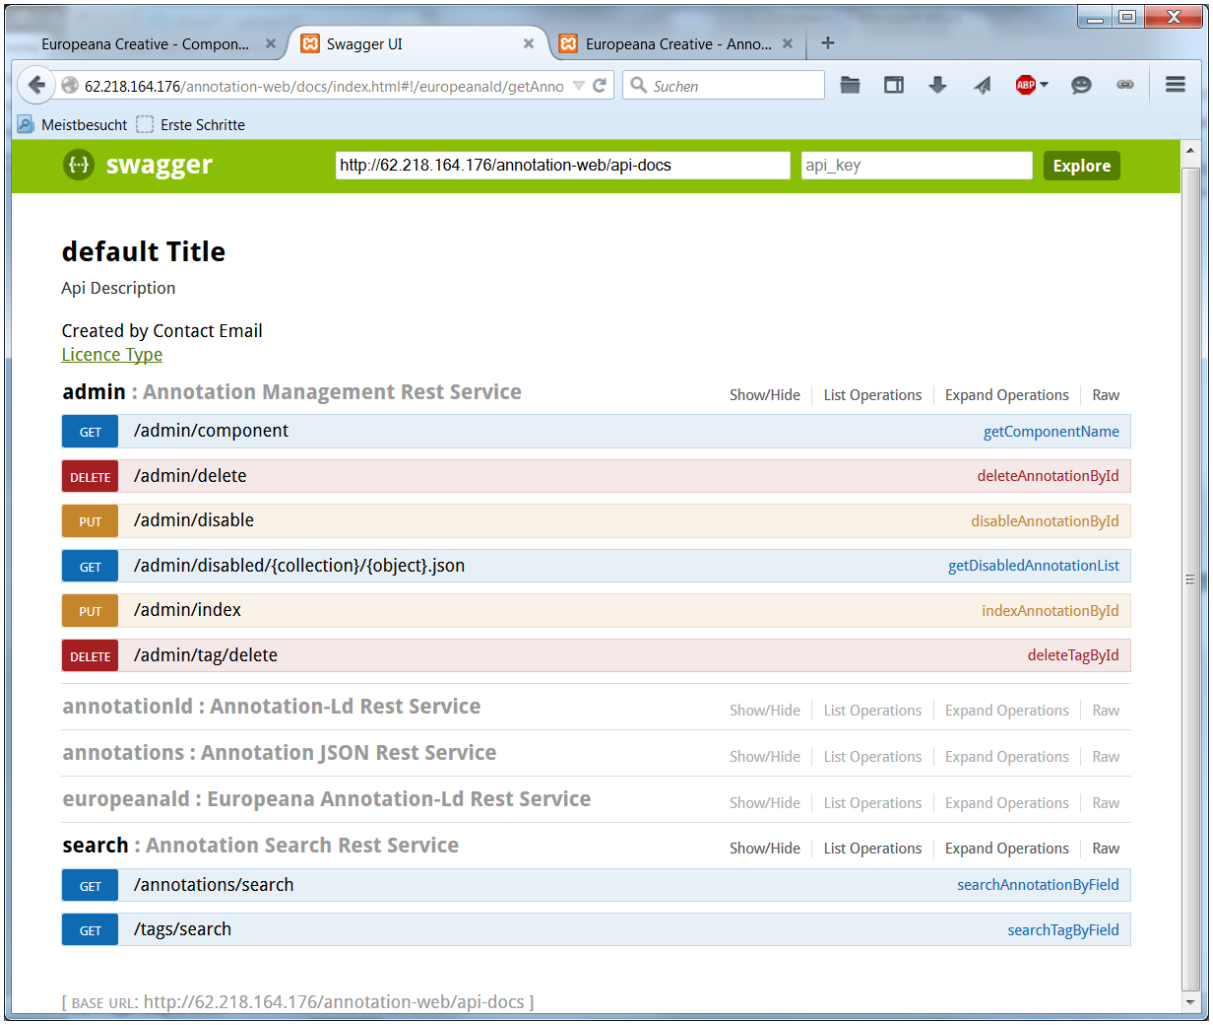

Image 4: Annotations Console - search and administration

1.3 User annotations in the Europeana Channels

Within WP4 Channel Development the Annotations API will be exploited to enable logged-in users of Europeana Portal Channels to annotate records and link items in different ways. The types of annotations with the highest priority will be tags, object interlinking and image annotation. Part of the development plan is also to enable users to create their own sets (small collections) of records, give them a title and a description and publish them for others to see. While users can choose to keep their annotations and sets private, it is the intention that users will mostly choose to make them publicly visible, and thus also included the Europeana search index to improve discovery.

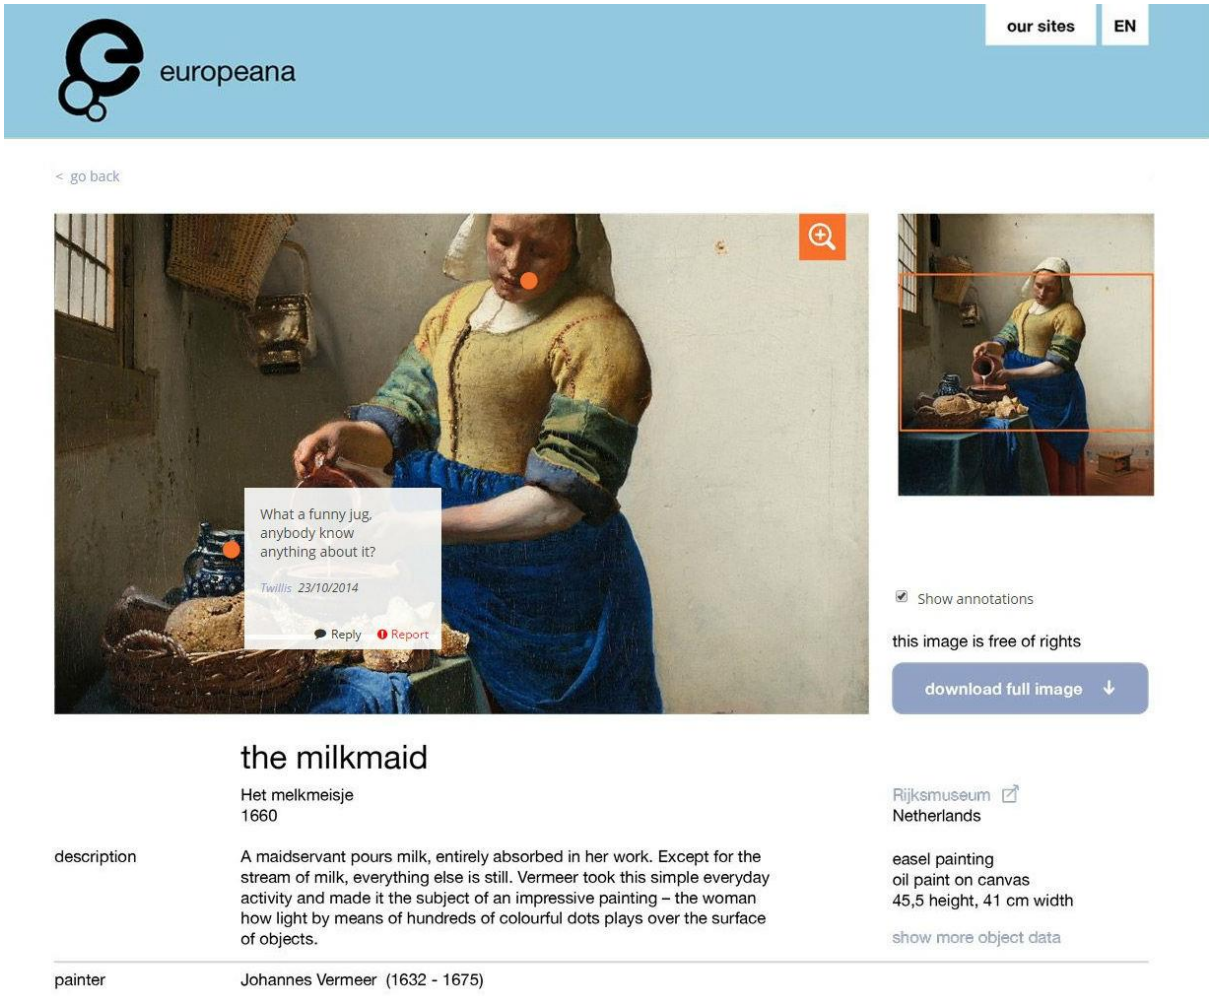

Image 5: Wireframe displaying the planned image annotation feature

The addition of user annotations and user created sets will require the development of a moderation framework to ensure that added annotations do not break the Europeana Terms for Use Contributions[14]. As we expect thousands of user annotations to be added (and in multiple languages) the moderation framework that is developed must be scalable. While technological measures form one part of the moderation framework, the most important part is "soft": active outreach to develop a sense of community among Europeana users and a feeling that their contributions are valued. This will encourage annotations to be made in a constructive spirit and for the community of users to largely moderate themselves, with intervention from Europeana and Europeana Channels partners only being a measure of last resort.

[14] See: http://www.europeana.eu/portal/rights/terms-for-user-contributions.html

Image 6: Discussing the Music Channel crowdsourcing features with the project partners in Athens

Development of user annotation features in the Portal Channels and specifically the Music Channel will be in focus during 2016. The specifics of some of these (crowdsourcing) features were also discussed during the WP2 workshop at the Training Days in Athens end of June 2015, and will be described around the time of the first beta release of Channels (July 2015).

1.4 WITH modules as addition to the system architecture

WITH is a platform that enables users to search across many sources of data and create annotations (i.e. link, enrich, group items) for the Europeana items. The annotations are represented in the Open Annotation format (using Europeana's guidelines[15] for the representation of annotations). Similar to the Historypin scenario, WITH is for Europeana and Europeana Sounds one additional front-end client where annotations can be created and stored in the Annotations API. WITH targets the partners of this project (Data Providers) and creative industries that want to re-use the Europeana Sounds content.

The WITH platform is an addition to the original plan made in WP2 the first year. Parts of WITH have been developed in the framework of Europeana Space (Ref 4) and EUscreenXL (Ref 5). This fact will allow us to reach a wider range of users and communities participating in these projects (creative

[15] These guidelines will be made available publicly at a later point.

industries, museums, artists, publishers and TV/broadcasting archives), bringing WP2 closer in achieving the critical KPI goal, which is the number of annotations (tags) added by users (KPI 9). Below is illustrated a use-case, using WITH where users can discover and annotate Europeana items and store them in the Annotations API.

1.4.1 Search

WITH enables searching across a variety of aggregators that provide access to cultural data (and content) through APIs. Europeana is the largest and most important source of data, but in addition to Europeana we bring data from Digital Public Library of America[16], National Library of Australia[17] and Digital NZ[18]. The purpose of this federated search is to also enable the linking of Europeana items with items from these external sources. In the rate of one per month, we are adding new external sources to the WITH platform (Rijksmuseum is scheduled for the current month). The figure below illustrates how users can search and get results from the available sources.

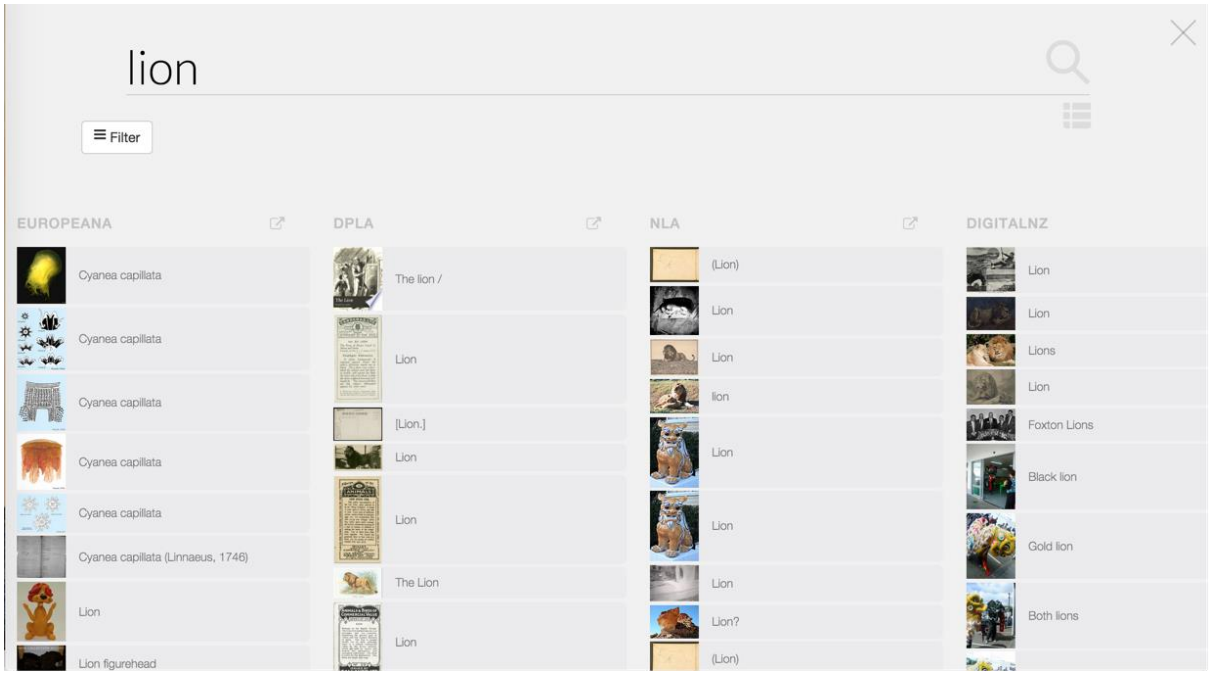

Image 7: Screenshot of the search functionality

Filtering: User can filter the sources of data (can choose which source to search and view) and can also filter the data using the filters provided by the respective APIs. In the figure below, the available filters that can be used to refine the search query are illustrated:

[16] See: http://dp.la/
[17] See: https://www.nla.gov.au/
[18] See: http://www.digitalnz.org/

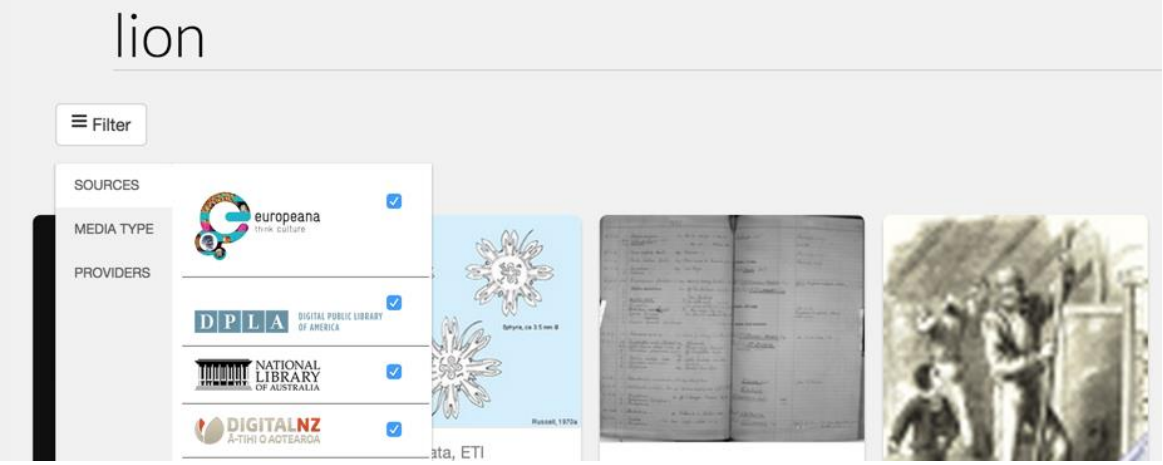

Image 8: Two filters to refine the search

Autocomplete: Enables users to quickly find and select from a pre-populated list of values as they type, leveraging searching and filtering. The Autocomplete widget provides suggestions while you type into the field.

Types of views: Infinite browsing of the data across the sources is facilitated with the Mosaic view. This functionality is illustrated in the following figure:

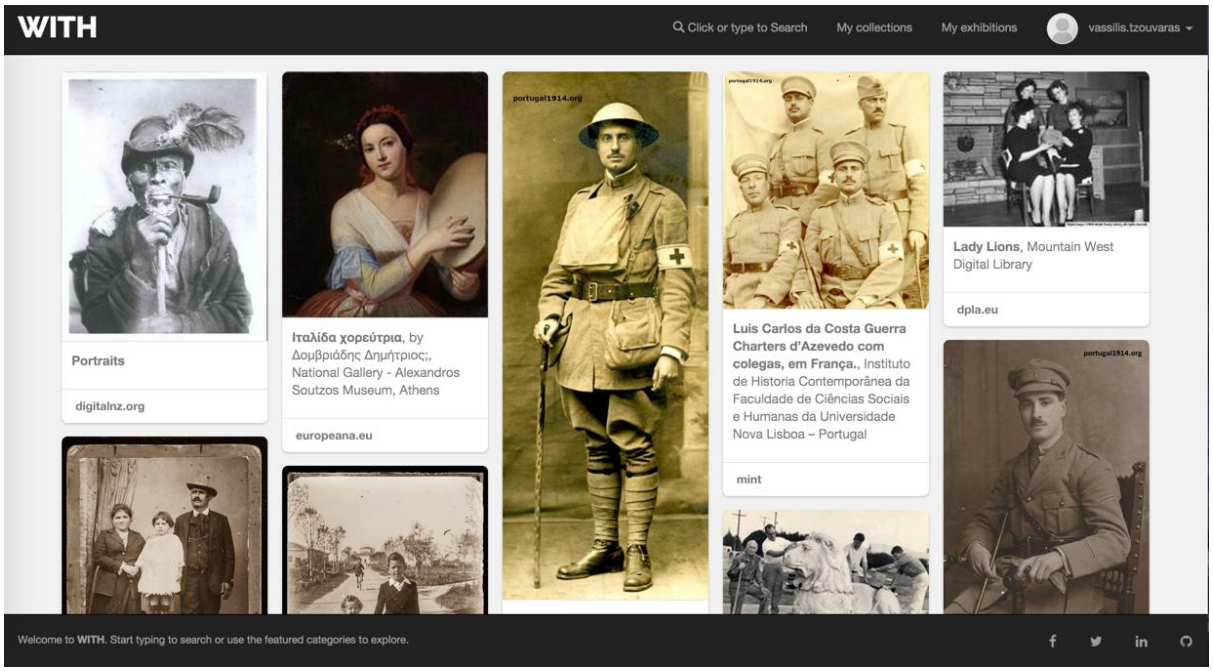

Image 9: Infinite browsing using the Mosaic view functionality

1.4.2 Create annotations

The aim of WITH is to enable users to create annotations for Europeana items. It currently covers three types of annotation; linking, enrichment and items grouping. The resulting annotations will be stored in the Annotations API.

Item Group: Users can create groups of items by clicking in the "collection" button. The user can then also add metadata (title, description, rights status, etc.) to these groups that contextualise and transform the group of items into a cultural collection. During development currently existing specifications for user created groups in the Europeana environment[19] will be taken into account.

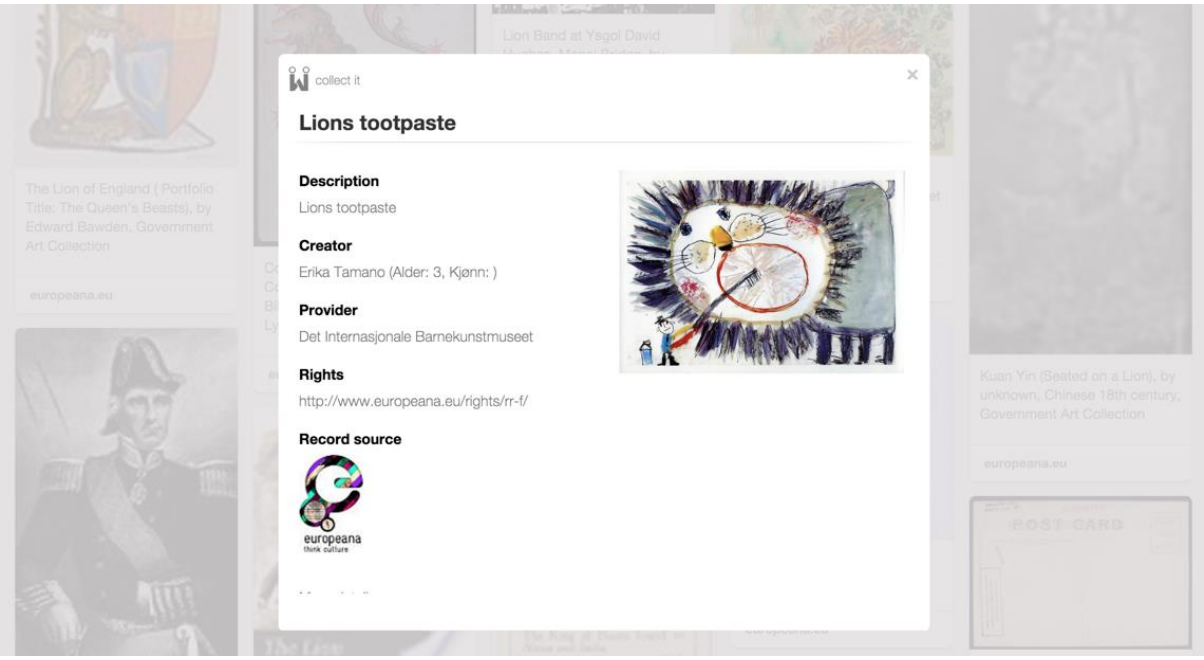

Image 10: Item view and collect item screenshot

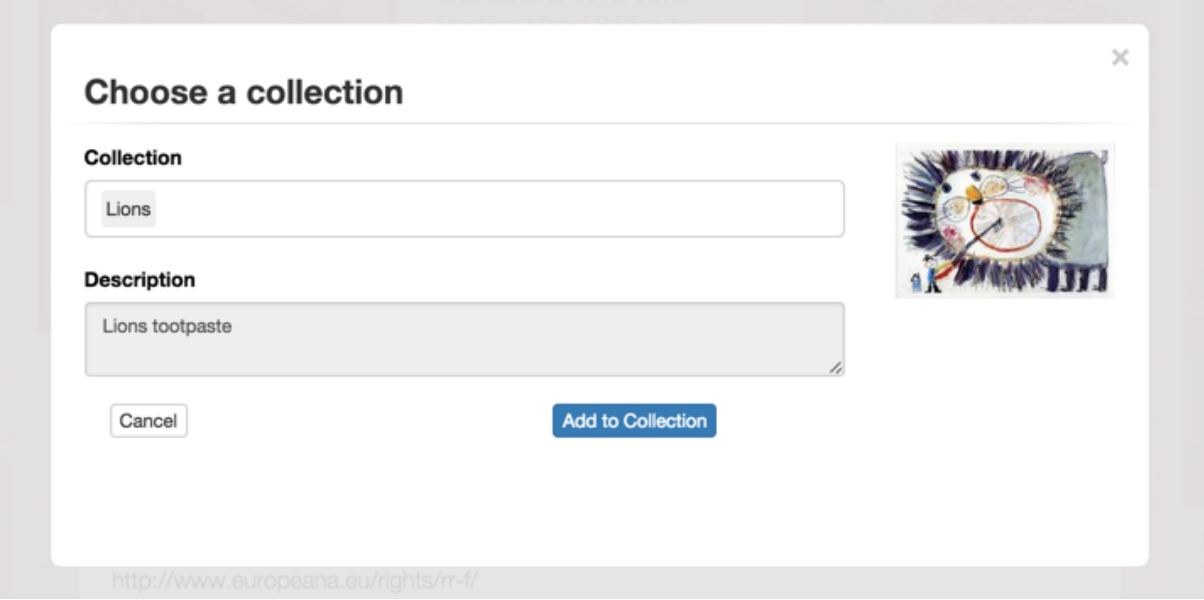

Image 11: Create a collection and add items in it

Manage groups: Logged-in users can manage their collections. They can add, delete make public or private a collection. Finally, users can export the collection(s) using the available API calls (illustrated

[19] See: https://docs.google.com/document/d/15nqqs7M9V25iku9NsiEfXl-vJpWh31LmyaYVA__cvho/edit#heading=h.gt4i94yyh7qh

in Image 13).

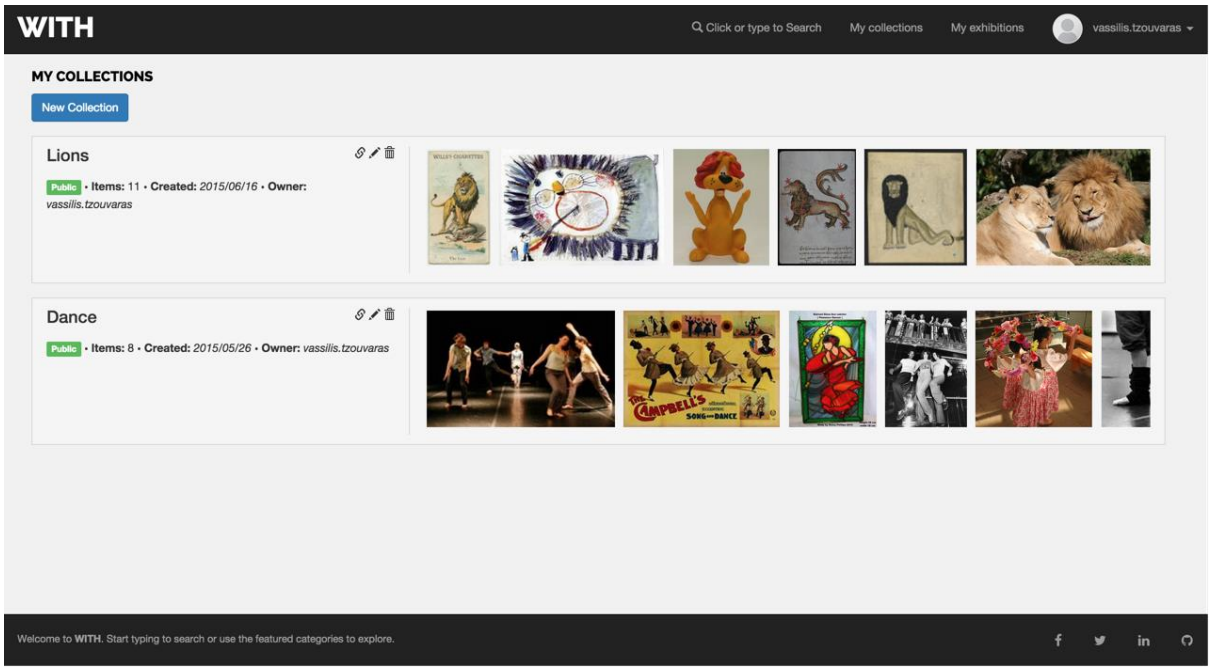

Image 12: Collection management module

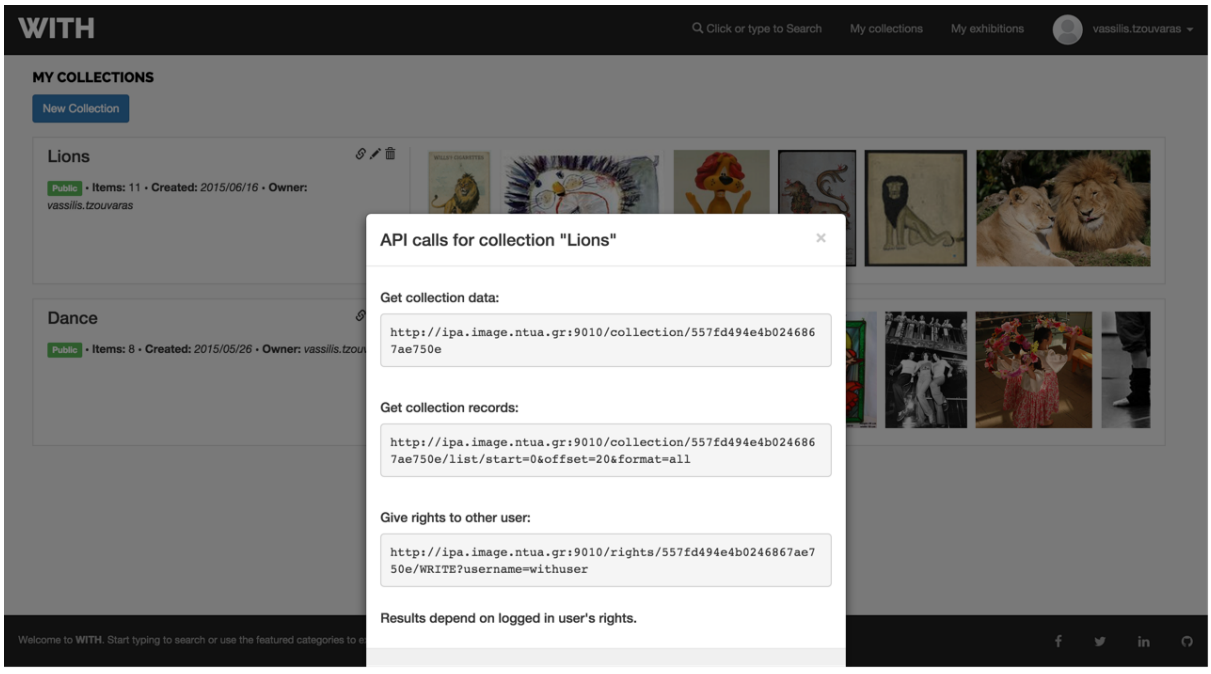

Image 13: API calls to retrieve collection and/or items

Linking and Enrichment

Linking and enrichment of items is performed using the Pundit tool form NET7. Using simple drag and drop operations, users will be able to link and enrich items using the Pundit functionality. This allows for further utilization of the Pundit tool and its specific functionality for enriching sound-related material within Europeana Sounds, next to the tool being available as a widget for the collection websites of Data Providers.

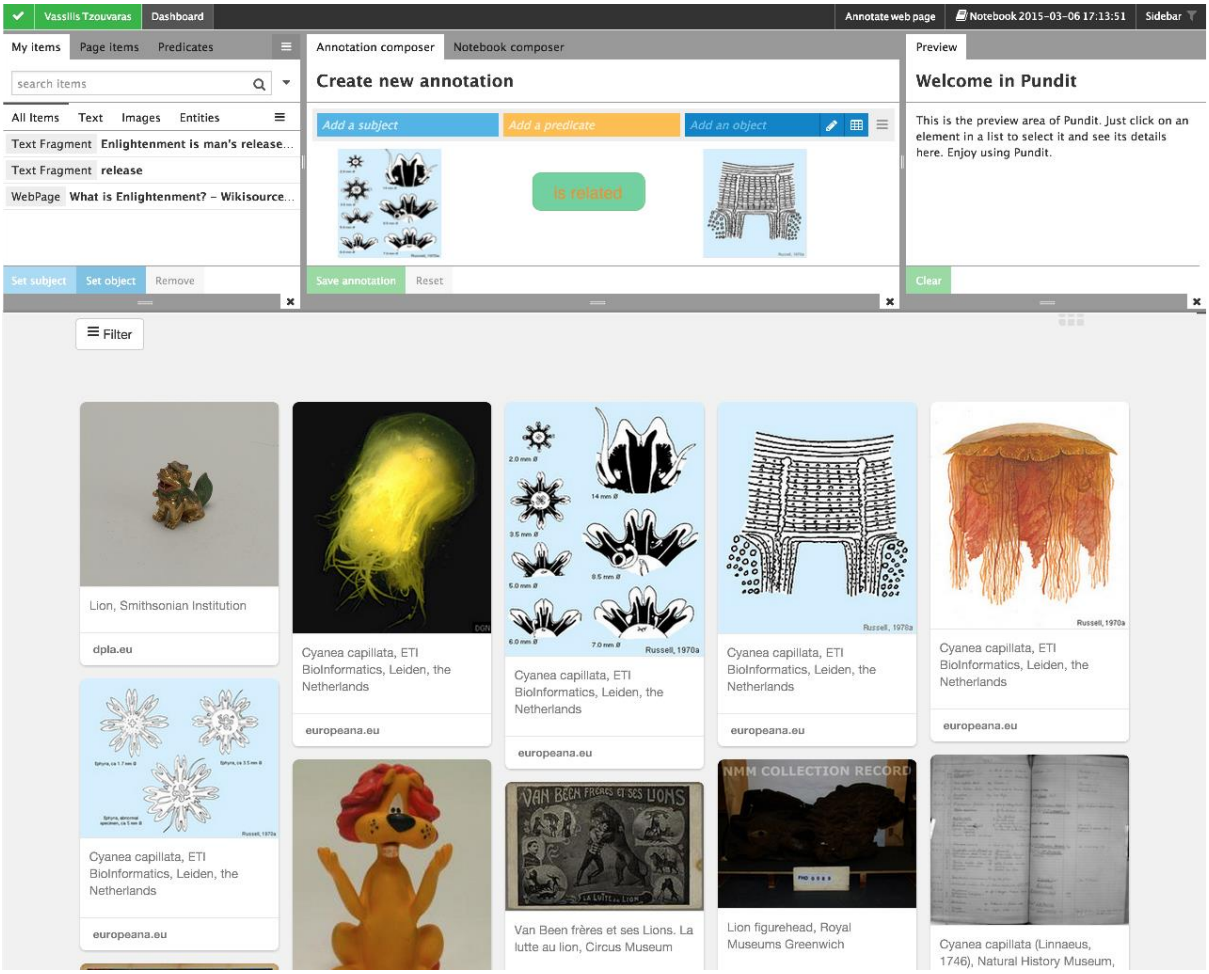

Image 14: Wireframe of the Pundit tool integrated in WITH

1.4.3 WITH export

WITH provides an API that can be used to export the user annotations in the appropriate format for ingestion in the Annotations API. The WITH API also gives access to all WITH services for usage by external services.

1.4.4 Current status of WITH and future development in Europeana Sounds

WITH currently support the basic functionality, which includes; services of the authentication, register/login, federated search, collection creation, collection management, and an API that; a)

gives access to these services and b) exports the annotations to the Annotations API.

In the framework of WP2 NTUA will develop the annotation services, including the integration of the Pundit tool. These improvements will allow users to create a link between Europeana items and links of Europeana items with external knowledge sources (LOD, authority files, thesauri, etc.). It is planned to have the first version of the WITH modules in October 2015.

During the WP2 workshop at the Training Days in Athens end of June 2015 some initial requirements were gathered together with the project partners.

Image 15: Discussing the WITH crowdsourcing features with the project partners in Athens

2 Historypin traditional music pilot

This second section describes the advancement of Shift with the *Crowdsourcing through specialised platforms scenario* (as introduced in D2.2), which focusses on the challenge of music identification. It describes the underlying user research, the proposed technical solution, first progress with the technical development, and a timeline for the development and integration.

As part of WP2, Shift is contributing to Task 2.1 - Crowdsourcing. As the crowdsourcing in WP2 is focused on micro-tasks that generate genuinely useful additions, Shift has proposed a traditional music pilot that brings archival sounds holdings to knowledgeable users that might not be aware of these collections by themselves.

2.1 User research

Anecdotal evidence gathered by Shift through informal conversations with traditional musicians suggested that a service to help traditional musicians identify and learn new tunes through listening to archival recordings would be useful. To understand this better, Shift undertook user consultation in the form of a survey to answer questions about how traditional musicians hear new tunes, how they identify and find them, and how they subsequently learn these new tunes.

The survey was open through March 2015 and attracted 59 responses in total (please refer to APPENDIX I for more details on the survey and users research by Shift). The traditional musicians who responded came from a variety of places, with the majority coming from the USA and Canada. Many of them play a range of instruments, and most of them attend local sessions.

2.1.1 How traditional musicians identify and learn new tunes

When asked how they heard new traditional tunes, the most common response was “from other players”. However the location varied - they might have heard it live at a session or a feis (a traditional music festival), live on the radio or recorded on a CDs YouTube clip.

A common workflow to learn a new tune was to record the tune they heard, find a similar ‘accepted’ or ‘canonical’ recording of that tune and then learning it by ear.

To do this, the musician would:

1. Record the tune at a session or a feis (traditional music festival)
2. Upload it into Tunepal it to find a name
3. Use the name to search for recordings of that tune on Google, YouTube or TheSession.org
4. Pick a recording they like best and learn that by ear

2.1.2 The need

This survey supported the anecdotal evidence gathered by Shift that there is a way to use the traditional music archive recordings that are being shared as part of Europeana Sounds to streamline

a process that traditional musicians are using in the real world as part of their daily engagement with traditional music.

The traditional music pilot lead by Shift is proposing to combine steps 2 and 3 in the above workflow and eliminate the need for musicians to first find the name of a tune and then use different sources to find high-quality archival recordings of it. As a result, musicians would find it easier to identify canonical tunes that they can learn by ear and to discover and connect to the traditional music archives that hold these tunes.

2.2 The solution: a Tunepal widget

Shift will create an app with a simple interface that utilises pre-existing tune identification software called Tunepal[20]. Tunepal is a technology that allows musicians to instantly obtain notation for a traditional tune based on live audio sampling of a performance. Available online and via mobile apps, Tunepal has over 20,000 users in 40 countries who make in excess of 1,000 music searches per day. Musicians can play a tune into Tunepal, which will then present the most likely name of that tune back to the user.

Image 16: Three screenshots from the Tunepal video showing the current functionality of the app on a mobile phone (recording, matching, and playback of the tune)

Shift is proposing the solution will use the Tunepal functionality to link queries (in the form of tune name, tune ABC transcription or tune recording) to specific archival recordings within Europeana

[20] See: http://www.tunepal.org

and to integrate this on a number of different websites, including thesession.org[21] (a forum for traditional Irish and Scottish musicians), Historypin.org and potentially the websites of Comhaltas Ceoltóri Eireann, Irish Traditional Music Archives and Tobar an Dualchais, the traditional music archives that are part of Europeana Sounds.

This app will take the form of a HTML/CSS/JavaScript widget that:

1. Allows the user to record a 10 second piece of music
2. Connects to Tunepal to get some suggested high-probability tune names
3. Queries the Europeana Search API to get back archive recordings that are likely matches
4. Allows the user listen to some of those audio candidates
5. Records (as an enrichment in the Annotations API) the link between input and matched recording, if the user indicates that there is a match.

2.2.1 Prototype Tunepal widget

Shift is currently working with Bryan Duggan, the developer of the Tunepal software, and Chris Xue, one of his students, to develop a web service interface for Tunepal that the JavaScript Tunepal widget can connect to.

At the time of writing, the Javascript widget has been developed and the connection to the Europeana API is being created. Below are some rough wireframes that show the app's user journey and the way the app will surface the archival content.

The first step is that the user records a 10-second bit of a particular tune (or in the later versions of the app uploads a file from their computer or pastes a link to a YouTube or SoundCloud file). This starts the Tunepal tune identification, which then gives the user a list of candidate tune names in order of probability, with the most likely match preselected. The user can also type the tune name into a search box, which gives them an identical list of tunes with the most likely match preselected. The name of this preselected tune is then sent as a query to Europeana and other external sources like tune books.

The next step will show the recordings that the query to Europeana has returned, with a 'play' button which allows the user to listen to a segment or all of the recording. There are also ABC-notations and sheet music for the user to explore. The user can then 'rate' these results based on whether they are recordings or sheet music of the tune they were looking for.

Once the user clicks 'Done', they get the option to copy and paste some HTML that allows them to share their results with others. The matches they've made between the tune name and the tune recordings will be sent to the Annotations API as contextual annotations.

[21] See: http://www.thesession.org

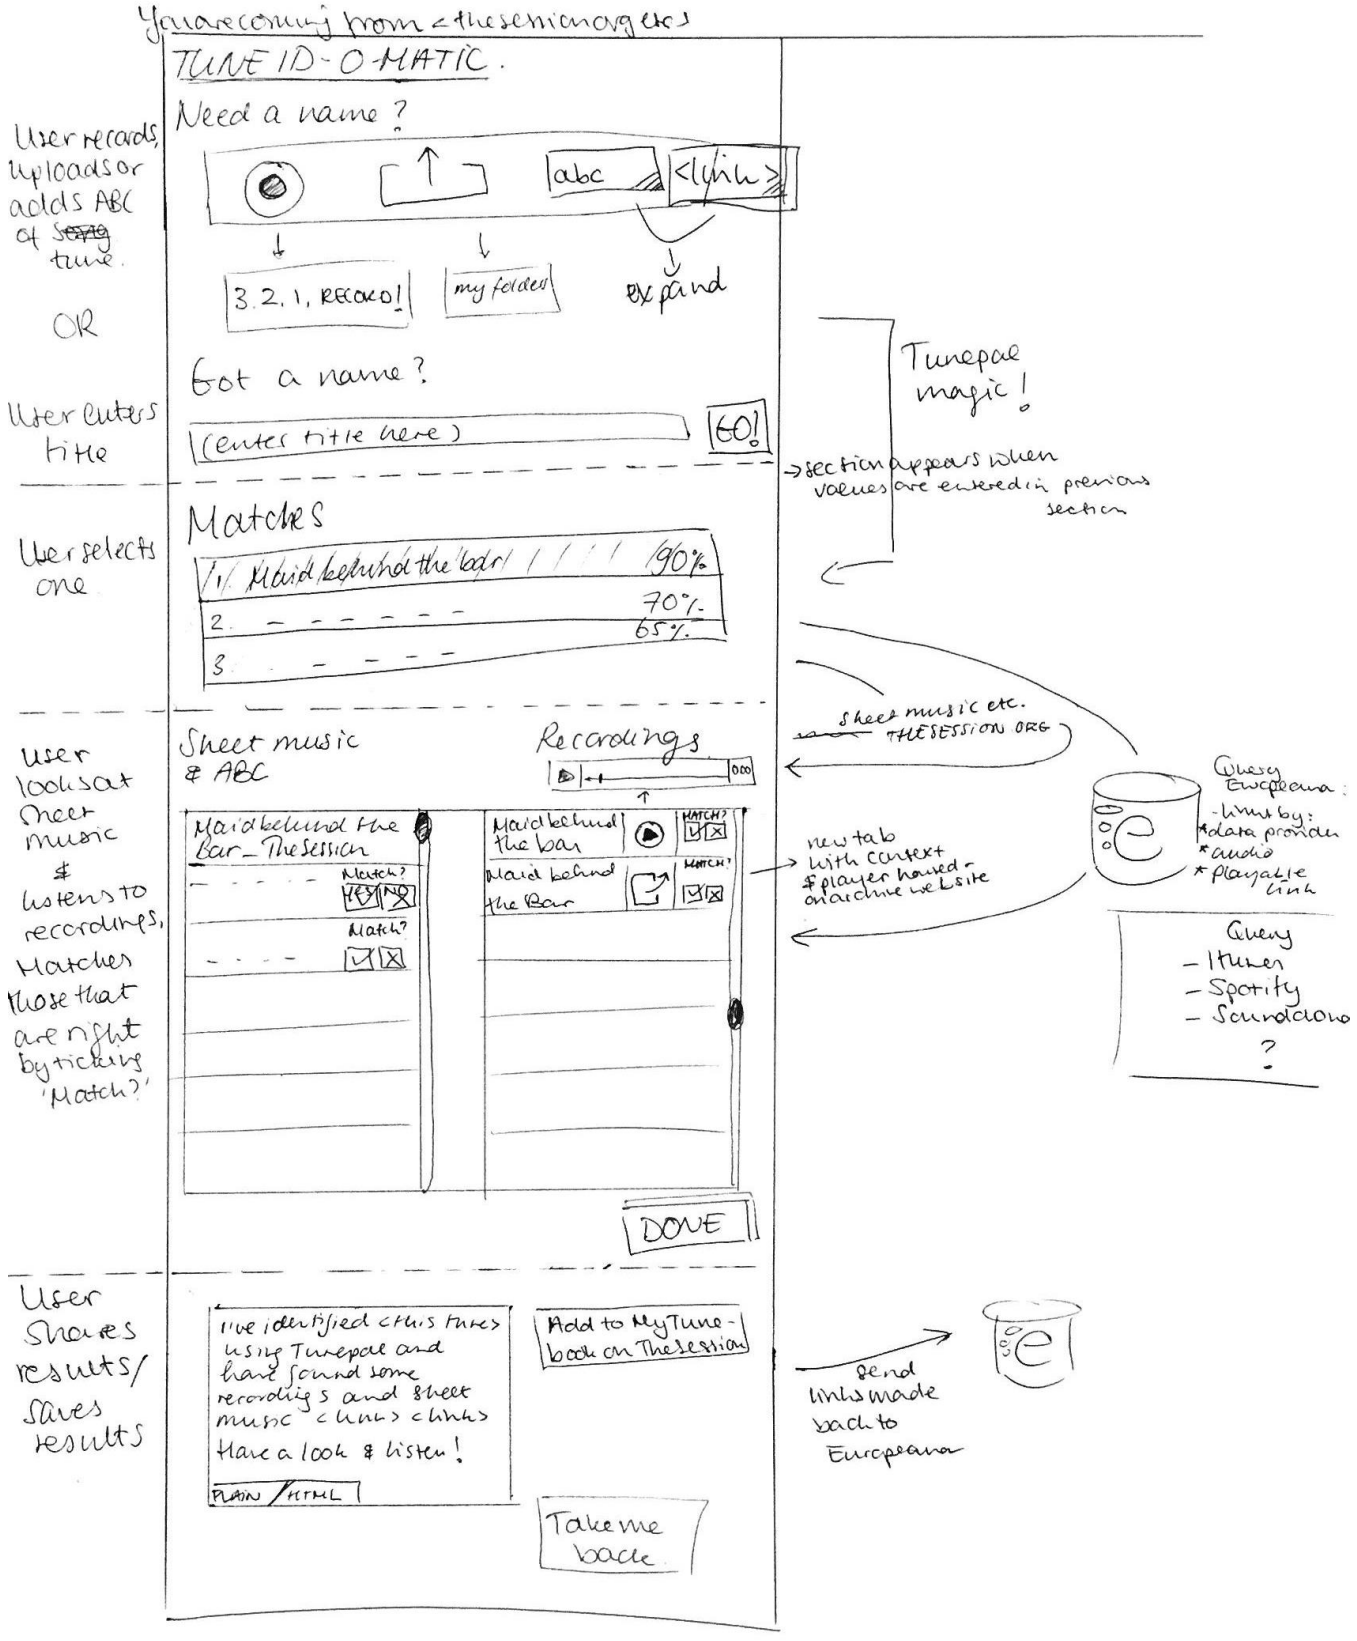

Image 17: A first sketch of the proposed traditional music app workflow

The first iteration of the app includes the 10-second recording functionality as sketched above. Shift will test this first version with a group of traditional Irish musicians to see how well it meets their needs. Together with Comhaltas and the developer of the Tunepal app, Shift will user test the functionality of the app at the Sligo Fleadh[22] a traditional music festival in August 2015.The app will be able to run in a web browser on a tablet, which will then be taken to the locations where the Scoil Éigse, the annual summer school, is taking place. The app can then be tested by a selection of the roughly 700 students and 100 tutors that are present. The data gathered from this user testing will be analysed and used to improve the app.

In addition, Shift staff will be doing some more general user research during the Fleadh, focusing on the needs and wants of traditional musicians and the way they are using digital tools to build up an archive or their local sessions.

Timeline:

May - June 2015	Testing of developed web-based Tunepal app
July - August 2015	Connecting the Tunepal widget to the Europeana API
10-16 August 2015	User testing the first iteration of the app at the Sligo Fleadh
September 2015	Evaluation of test at the Fleadh

2.2.2 Future iterations

Shift will incorporate the feedback from the user testing and iterate the prototype. The next version of the widget will be an online version on TheSession.org, a forum for traditional musicians and a final version will be integrated into Historypin.org early 2016. This will open up access to the traditional music archives in a place that traditional musicians are already visiting and will help test the widget in an online environment.

Shift will create crowdsourcing tools for traditional musicians on the Historypin.org website. These tools will incorporate ways to share photos, audio recordings and videos, as well as ways of annotating these.

Additionally, there is potential scope to extend the widget so it can use pre-recorded tunes (for example from a session in a pub) as input (see sketch above). This will require additional functionalities like an audio cropping or selecting tool to select the individual tunes in a set. This extended widget can then become part of the crowdsourcing tools on Historypin.org

[22] See: http://www.fleadhcheoil.ie/

Timeline:

September - December 2015	Deploying Tunepal widget on thesession.org for testing and scoping extensions to the widget. Creating Historypin for Trad crowdsourcing tools
January - April 2016	Developing Historypin for Trad further and running engagement events with traditional musicians to generate contributions
May - July 2016	Evaluation of crowdsourcing tools and contributions

2.2.3 Relevant general Tunepal development information

A REST API is being developed for Tunepal using Java with Jersey[23]. This will make it possible to use Tunepal as a web service and to build new applications for the core Tunepal archive and matching algorithms.

A new responsive web app is being developed using Javascript, HTML5 and Angular.js[24]. This app will run on all modern browsers, including mobile browsers, and allow users to make audio queries straight from the browser. Additionally the web app will display content from Europeana Sounds by making API calls against the Europeana API's. The web app is being developed to be embeddable in other websites.

All the Tunepal source code and API's will be open sourced and available on a public GitHub repository[25].

2.3 Lessons-learned and recommendations

In conclusion, Shift is aiming to help traditional musicians engage with archive collections containing traditional music more frequently and more deeply, by connecting websites popular with traditional musicians to the collections of multiple archives. By linking archives and traditional musicians using digital tools, we can reduce barriers in teaching and learning. This also maximises the possibility for reuse of the metadata aggregated within the Europeana Sounds project, and makes discovery and reuse of archival materials more meaningful for specific knowledge communities.

By engaging with the active traditional music community through surveys and social media, Shift has been able to define a particular need, i.e. a quicker way of identifying tunes and a more intuitive way of exploring traditional music archives.

[23] See: https://jersey.java.net/
[24] See: https://angularjs.org/
[25] See: https://github.com/skooter500

As seen in Athens in June 2015, other Data Providers who do not have traditional Irish and Scottish archival material were interested to see how they could use the learning developed as part of the traditional music pilot. Extrapolating the learning from this pilot into general lessons learned about archives and their audiences, and sharing this with the Europeana Sounds project, will be a focus point for Shift in the second half of the project.

Image 18: Discussing the Tunepal crowdsourcing features with the project partners in Athens

By working together with the developer of an existing tool that already meets some of these needs, Shift is hoping to create a lasting legacy for the Europeana Sounds project, in the shape of an open-source app that uses the Europeana API to help users explore sound archives.

3 Enrichment annotation scenario pilot

This third section describes the advancement of NET7 with the *Manual correction of semantic annotations* scenario (as described in D2.2), which focusses on enhancing existing metadata records through a crowdsourcing widget based on Pundit, and the evaluation of the current state of this work. It describes the development of a pilot version of the crowdsourcing widget, user research that has been applied to this pilot version, describes the problem areas that emerged from the evaluation (including recommendations for improvement), and finally looks ahead at the upcoming work, including the next pilot scenario.

3.1 User research

In order to let NISV perform user research, Net7 created a pilot crowdsourcing widget, based on Pundit, for creating enrichments with information about sound-related media objects (for this pilot only audio objects were used), which is intended to be used on the website of Data Providers. The application is based on the metadata of audio resources already ingested by Europeana (from NISV), to ensure this pilot already fits with data flows as illustrated in Image 1 of this document. More information about the technical development will follow in section 3.2.

Image 19: The widget (top dark bar) as used during the user research

This pilot is effectively the first prototype of the Enrichment Annotation Scenario (EAS) or manual correction of semantic enrichment (as described in D2.2). This user research is executed to give an end-users' perspective on the widget. For this research people fitting the target audience for Europeana Sounds were selected (those who are interested in sound-related cultural objects, who can be either considered a Culture Snacker, or a Culture Vulture).

The main functionalities that were tested by end-users in this first user research of the widget, in a chronological order, were (Referencing the User Stories and Epics from D2.2, Appendix I):

1. Add a comment (free text) to the audio object (D2.2 Epic number E4, User Story S15)
2. Add a relevant music genre from a fixed vocabulary to the audio object (this is taken from the SKOS Genre vocabulary developed within the work of WP1, see D1.3: Ontologies for Sound, Ref 4) (D2.2 Epic number E1, User Story S29)
3. Link the audio object to a relevant digital object from Europeana (D2.2 Epic number E6, User Story S16)
4. Add relevant information from DBpedia and Freebase to the audio object (D2.2 Epic number E6, User Story S20)
5. Delete an user-created enrichment (D2.2 Epic number E5, User Story S10; implemented only for the user that is the creator of the annotation)

3.1.1 The current enrichment workflow in Pundit

In order to actually enrich an object using the Pundit widget, whether it is an added comment, a fixed concept from one of the repositories or creating a link between an object and another object from Europeana, the user always needs to perform several steps to properly save the annotation. Chapter 3.1.6 shows that some participants came close to creating an annotation, but that they sometimes lacked the ability to save it in the correct way. Underneath a quick walk-through is presented in order to show how the widget gives feedback in the steps for saving an annotation.

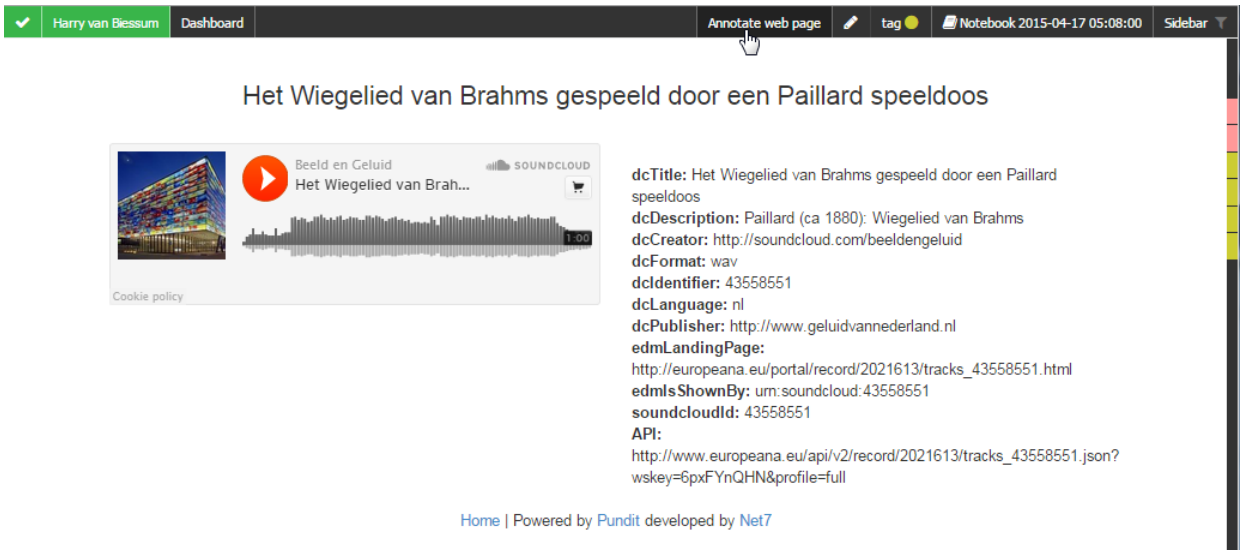

Image 20: After being logged in, the button "Annotate web page" needs to be clicked on in the top menu bar (Step 1).

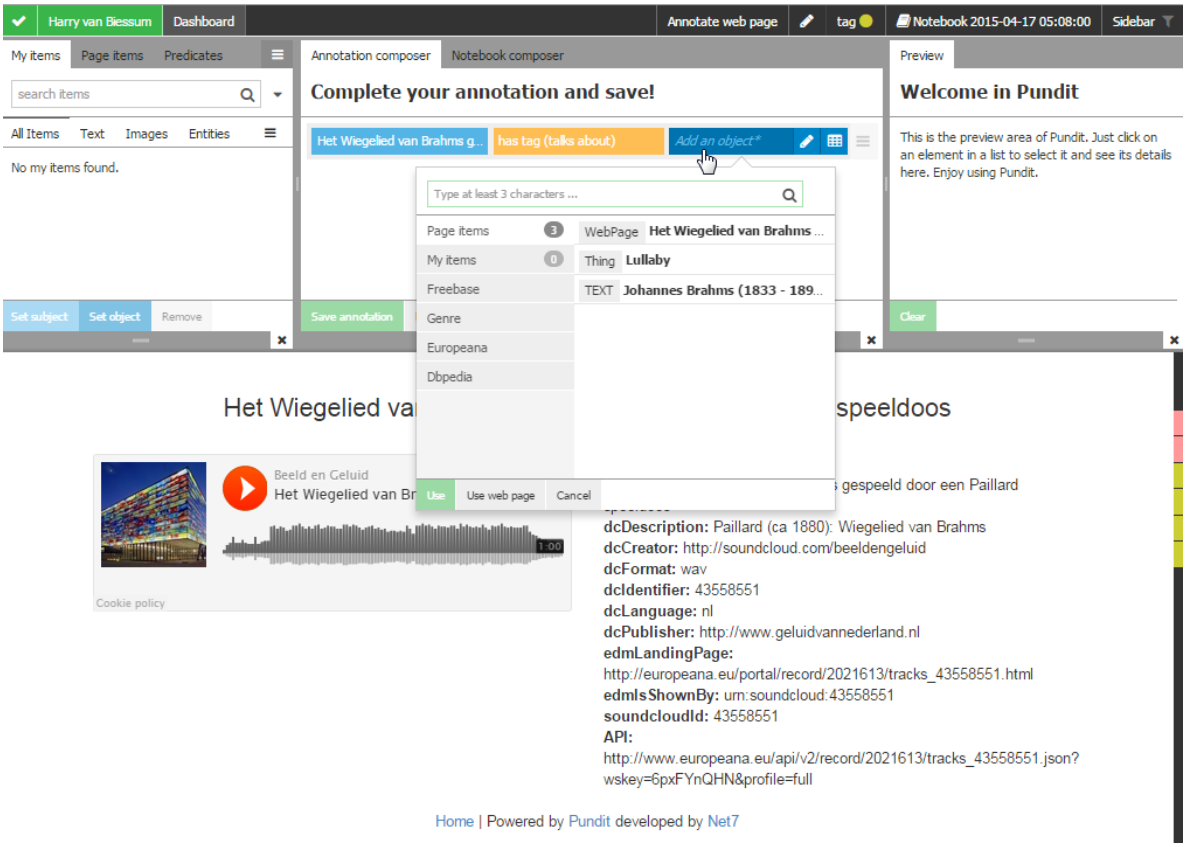

Image 21: The dashboard unfolds and the next step is to click on 'Add an object*', which causes the search field to unfold (Step 2).

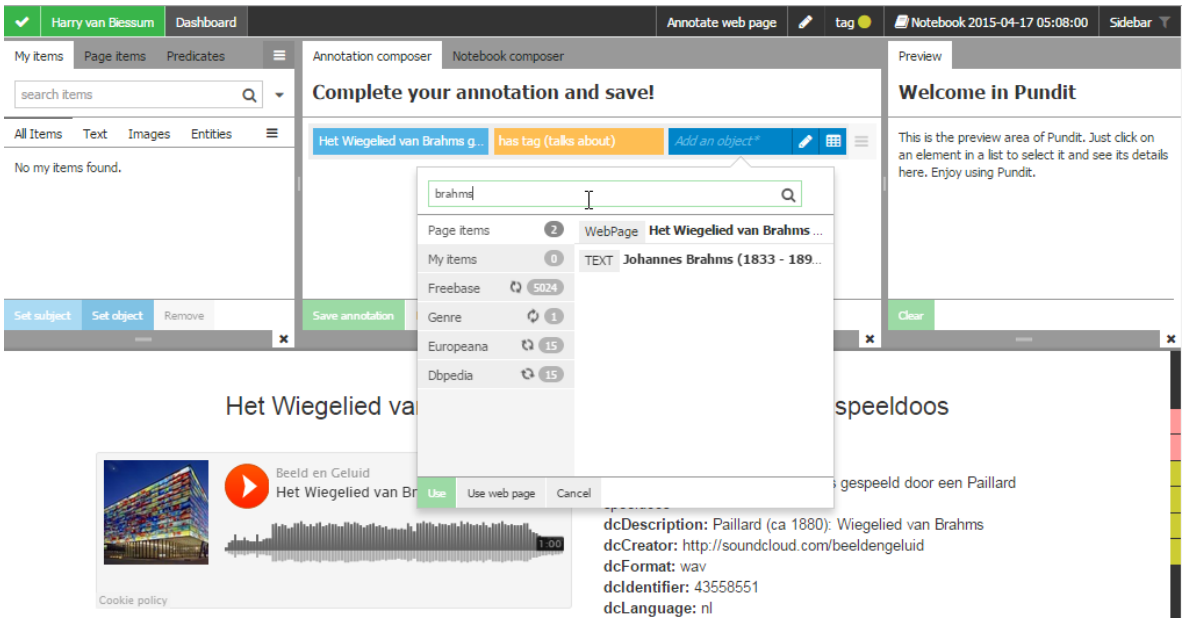

Image 22: Whilst typing "brahms" in the search field, the four different repositories are being checked with this query, in this process the symbol of two rotating arrows and a number next to it appear to indicate that the query is in progress and what the number of hits that are found based upon the query so far (Step 3).

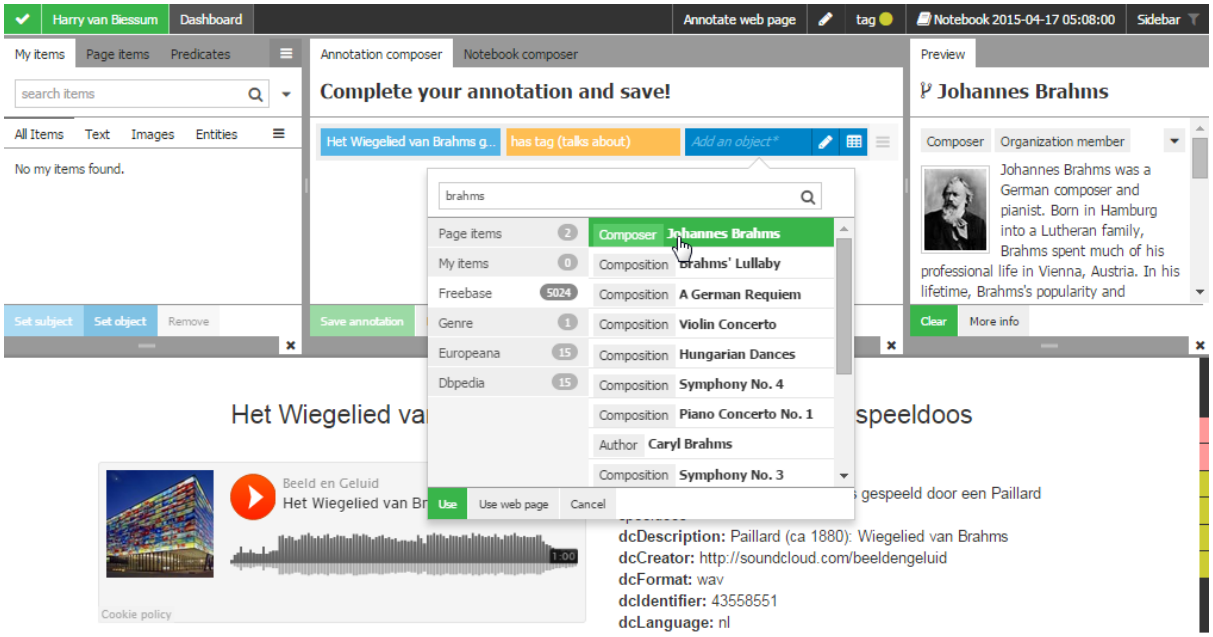

Image 23: After clicking one of the repositories (in this case 'Freebase') the results of the query are listed and when clicking on one of these results the Preview window on the right, the widget provides information about the concept, in this case the composer Johannes Brahms (Step 4).

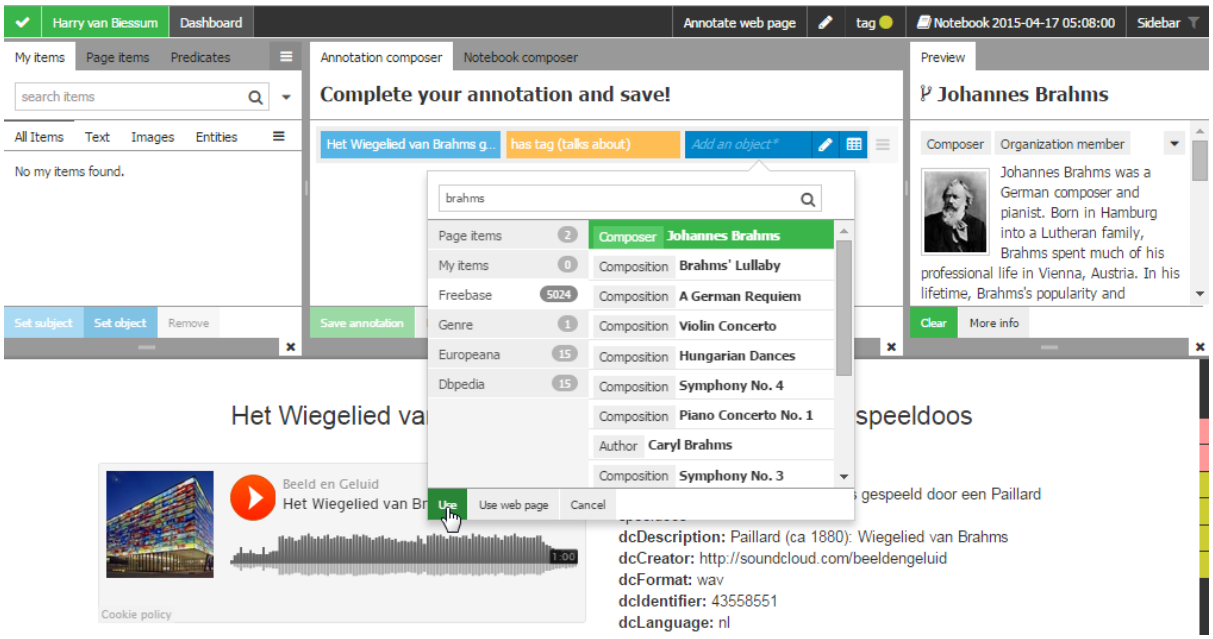

Image 24: After checking if this is indeed a concept that is relevant to enrich the object with, the green button "Use;" needs to be clicked on (Step 5).

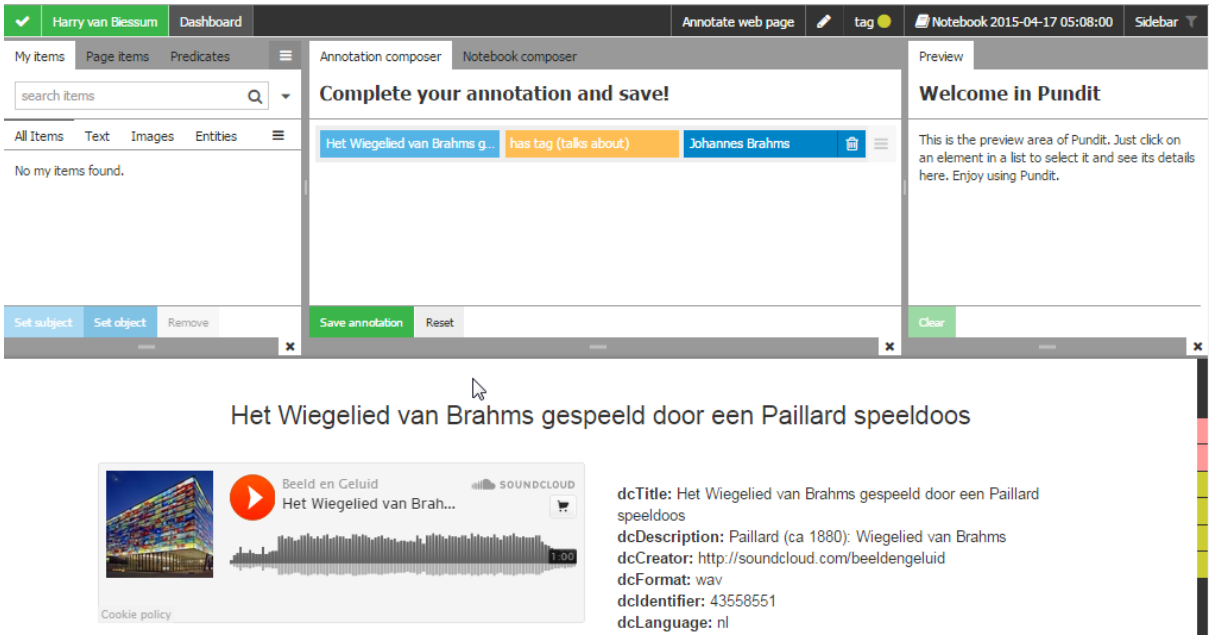

Image 25: After clicking 'Use' this screen appears, the chosen concept 'Johannes Brahms' is now written in the right-sided blue box. About half of the participants who came to this stadium in the user research thought the annotation was completed by now, which is not the case (see Image 36 in chapter 3.1.6 for more details).

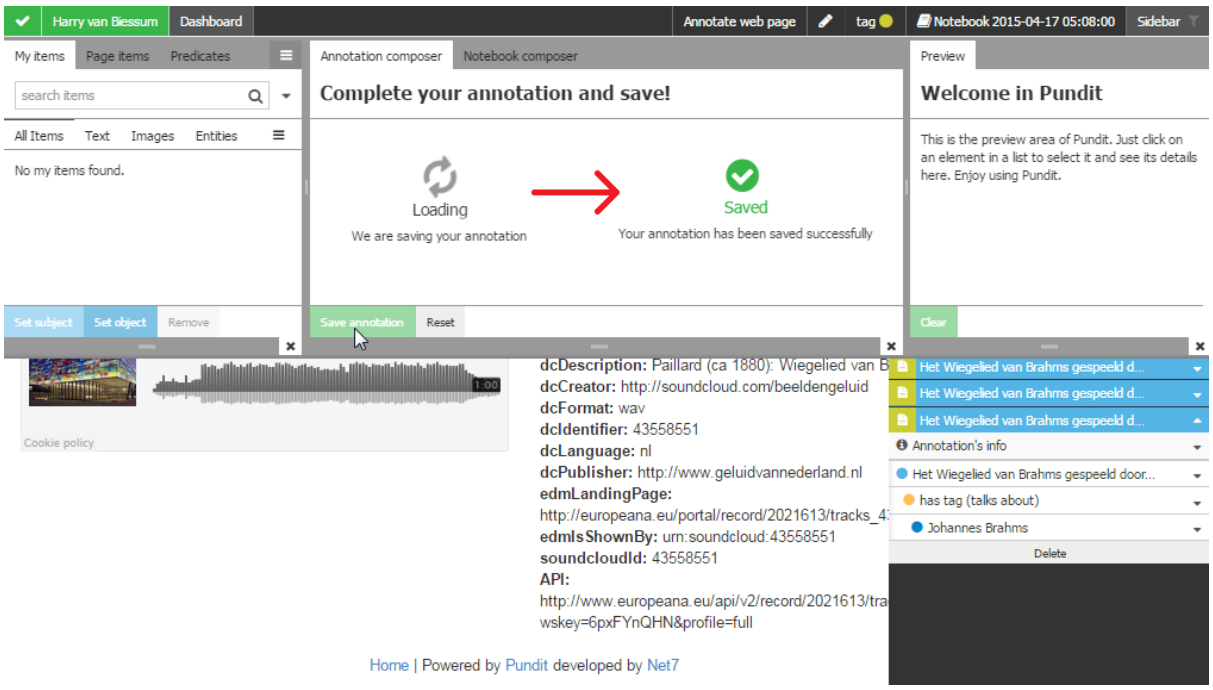

Image 26: To properly save the annotation the button "Save annotation" needs to be clicked on which causes the widget to immediately start giving feedback in the screen in the middle. First the text "Loading: We are saving your annotation" appears with an animation of two rotating arrows and after some seconds a green symbol appears with the text "Saved: Your annotation has been saved successfully". The sidebar (low-right) appears with all the annotations that were saved (from all users) and then the dashboard folds again, except for this sidebar (Step 6).

3.1.2 User evaluation of the current Pundit enrichment tool for sounds

For the qualitative user research participants came to NISV in Hilversum, and used the widget on a laptop (with an external mouse) for approximately one hour. The participants needed to complete several tasks, and during these tasks they were asked to tell as much as they could about what they expected from and thought about the widget. The tasks and questions were divided into 12 parts (please refer to APPENDIX II for all the questions and tasks, translated to English). In the beginning, the questions and tasks were about what the participants' expected before the goal of the widget was explained, and after all the tasks some general questions were asked. The participants also needed to give grades between 1 (very bad) and 10 (very good) for different aspects of the widget (plus a reason), namely on:

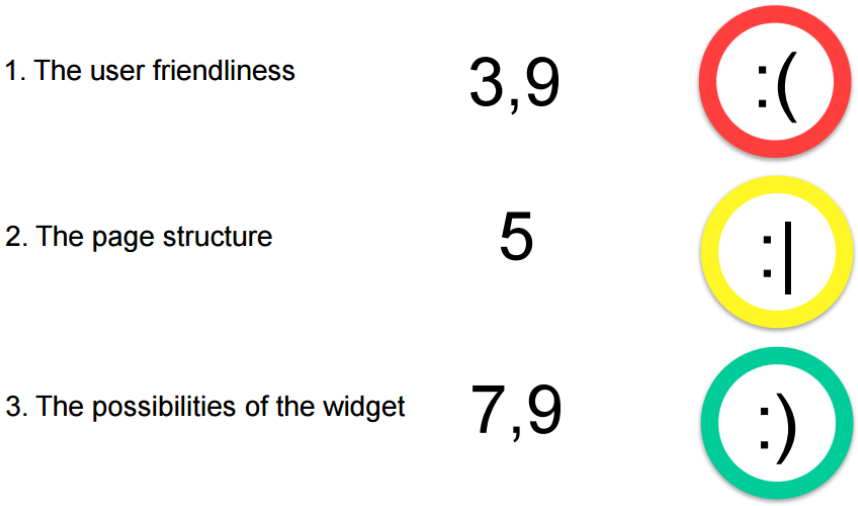

Image 27: Average grades for the widget given by the participants in the user research

In order to get enough participants for this user research social media was used. Facebook posts with an invitation on Muziekencylopedie.nl[26] and Geluid van Nederland[27] were sent out, and an invite[28] from the Twitter account of NISV was send to its 11.900 followers. Unfortunately only one person responded to all these messages and eventually came to Hilversum to test the widget. Three colleagues from other departments at Sound and Vision who have affinity with sounds, music and/or audio, but no relation to the Europeana Sounds project, also helped by taking part in this user research. Furthermore, the coordinator of the volunteers who work at the NISV (more than 70 at the moment) made a shortlist with people who professionally worked with audio, sounds and/or music projects and who are now retired. Four were able to give their perspective on the widget. The average age of the participants was 50 years and 10.5 months, with the youngest 23 and the oldest participant being 80 years old.

[26] See: https://www.facebook.com/pages/Muziekencyclopedienl/208878199150523?fref=ts
[27] See: https://www.facebook.com/GeluidvanNederland?fref=ts
[28] See: https://twitter.com/BeeldenGeluid/status/588660321377615872

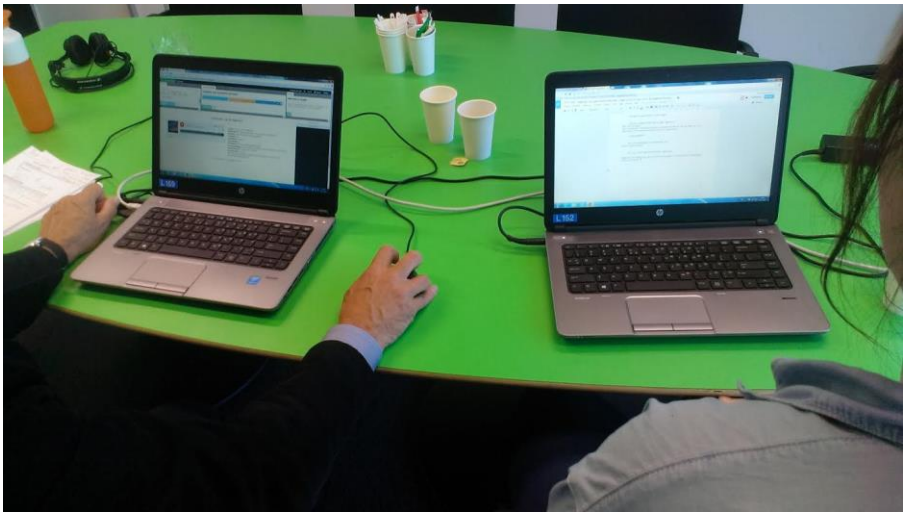

Image 28: One of the participants using the widget and a colleague taking notes (photo: Harry van Biessum - CC0)

All participants walked through most of the functionalities of the widget. The questions and tasks were divided into four different phases:

1. Expectations and exploration (questions and tasks from part 1 till part 5)
2. Enriching the audio object (questions and tasks from part 6 till part 10)
3. Delete annotation (task in part 11)
4. General questions about the widget (questions in part 12)

To gain a proper sense of what the participants thought the widget was about, the goal of the widget was not explained until phase 2. After that the participants were not told how they could execute the functionalities with this widget up until phase 3, in order to first see how they would use the widget intuitively, without any guidance. Just before phase 3 (task: delete an self-generated annotation) the explanation of how the widget actually works was given and in phase 4 several general questions about the widget and the research itself were asked (in this part the participants also needed to give the above-mentioned grades).

Several different problem areas emerged from the user research conducted by NISV based on the pilot crowdsourcing widget that was developed by NET7. All these areas will be explained in further detail in the paragraphs below, and examples are provided from cultural heritage crowdsourcing websites to show possible solutions for the different problem areas in the field. An important problem area is that most participants had the feeling that (1) the field in which they can search all the different repositories was hidden. There were also problems concerning the (2) feedback that was expected by the users after clicking buttons. The (3) wording of the interface elements, like buttons and fields names, is also a problem area for almost all users, besides most participants wanted the widget to be in Dutch (their native language). Another problem area is the (4) lack of guidance from the widget that the participants experienced when trying to enrich the audio object.

3.1.3 Search field for repositories is hidden

Before users can actually enrich an item with an entity from one of the repositories (DBpedia, Freebase, PURL or Europeana) using the widget, they need to find the search field that will search these databases with enrichment resources (targets). It takes at least two clicks on the item-page to get this search field in display after being logged in, as shown in steps 1 and 2 in the previous chapter (See: Image 20 and 21).

The general impression from the users is that this specific search field stays hidden in all the different menus they see. From the eight users that explored the widget and that were given several tasks to use one of these repositories, only two actually filled in their query in the proper field at the first attempt. Most got stuck in the first search field that is available after opening the Dashboard, which is in the search field in the menu “My items” that is inviting users to “search items”.

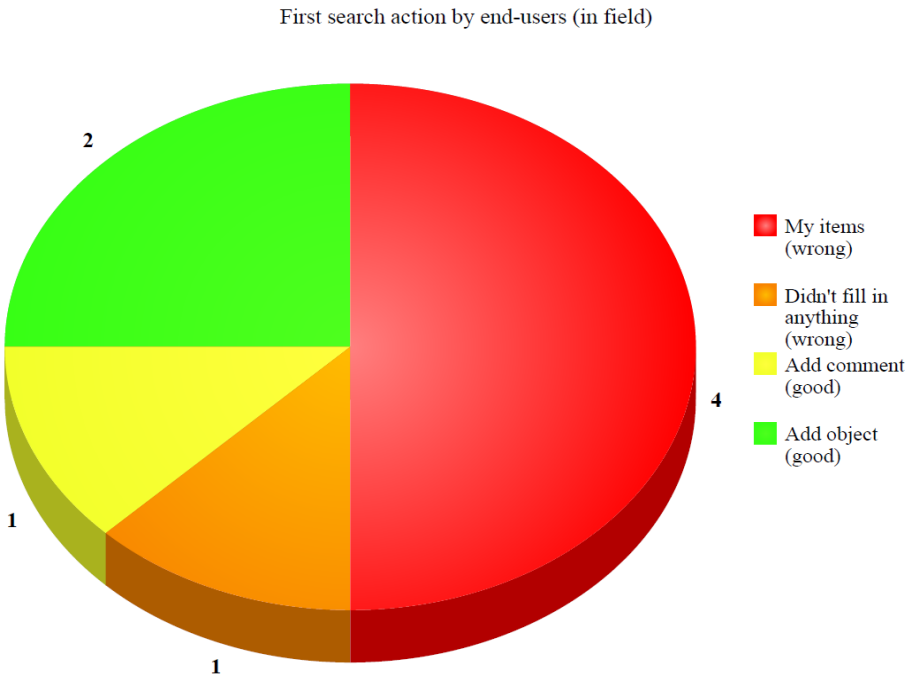

Image 29: First search action by the participants in different fields

The fact that most users could not find the proper search field resulted in a lot of frustration for the users and less enrichments during all the sessions. In chapter 3.1.6 the results of the number of participants that completed the different tasks properly, almost and not at all are shown. Overall, all tasks were not completed properly by almost all the users.

A recommendation to improve this problem area is to carefully consider what would be the first search field that users come across in the new version of the widget, and where this field would be (prominently) placed within the widget. The World Service Radio Archive (prototype), created by the Research & Development department of the BBC, displays their search field for fixed keywords

(corresponding to Wikipedia articles) immediately.[29] In this way possible contributors can directly add a tag without having to click two buttons before getting the right search field. Note that users with small screens possibly need to scroll down to arrive at this search field.

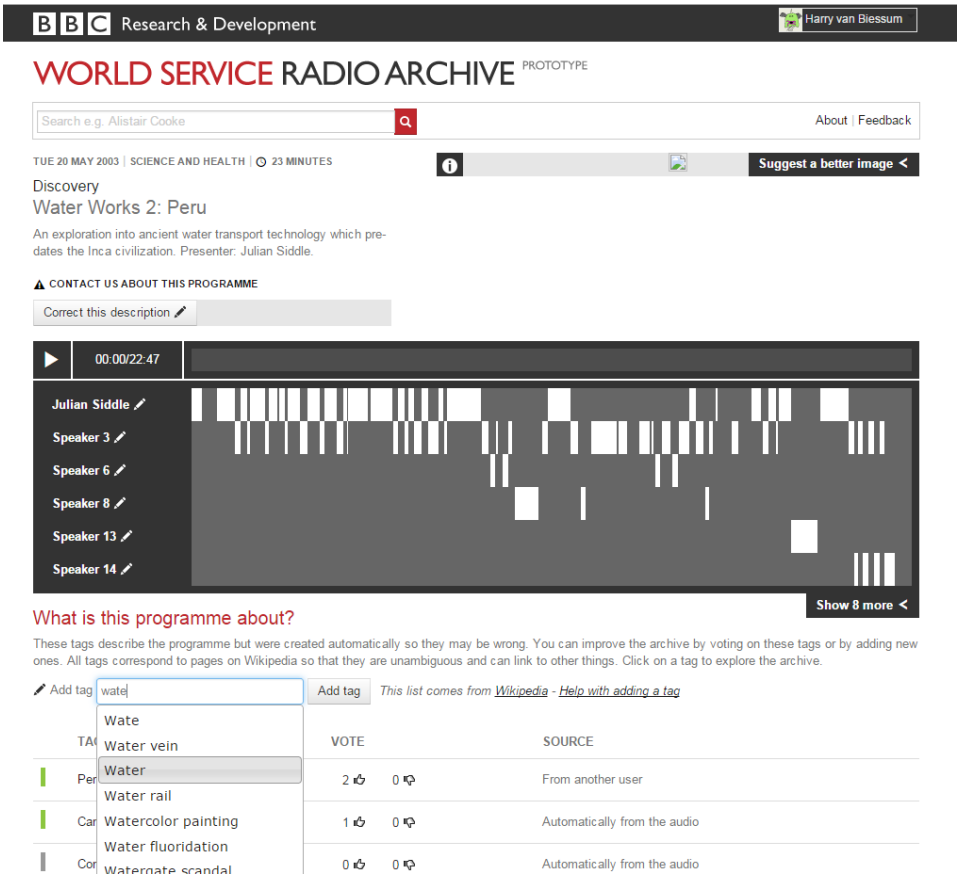

Image 30: Screenshot of the search field used on the platform World Service Radio Archive (prototype) for adding tags immediately.

SoundCloud[30] has a similar way of facilitating the possibility for users to enrich audio objects, but on their platform users can add comments instead of fixed keywords from a repository. Users are invited with the simple text 'Write a comment...' just underneath the audio player.

[29] See: http://worldservice.prototyping.bbc.co.uk/
[30] See: https://soundcloud.com/

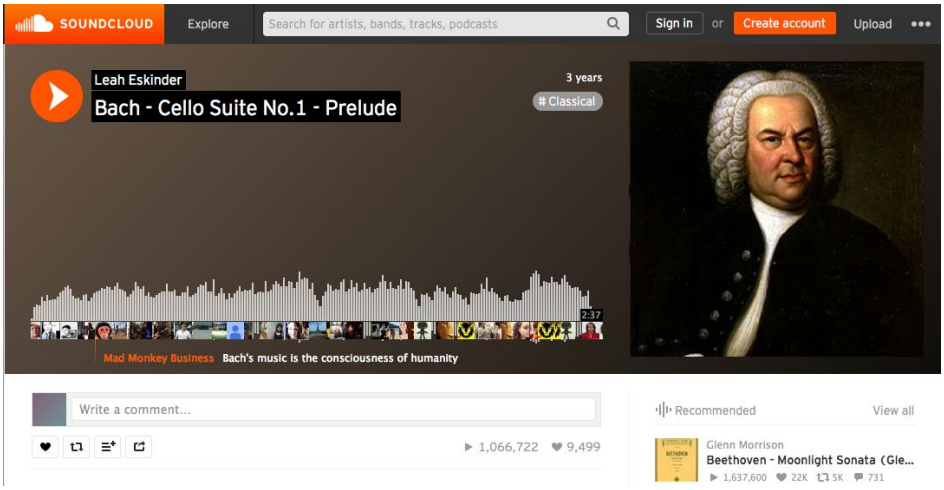

Image 31: Item-page on SoundCloud with the possibility to add a comment underneath the audio player.

3.1.4 Feedback

Most users who used the Pundit widget wondered what happened after they clicked on a button or searched for something in a field. It was not clear that they actually done something by clicking on a button. Most users lacked feedback from the interface, when an action was executed. For instance some users asked explicitly 'have I clicked something now?' when clicking "Annotate Web Page" (once the Dashboard was already open). As can be seen in the screenshots underneath the only visual feedback that is given after clicking 'Annotate Web Page' is that the font in the long shaped light blue box changes in colour and the title of the item appears there (see red arrow in Image 32 below, to see the changed box).

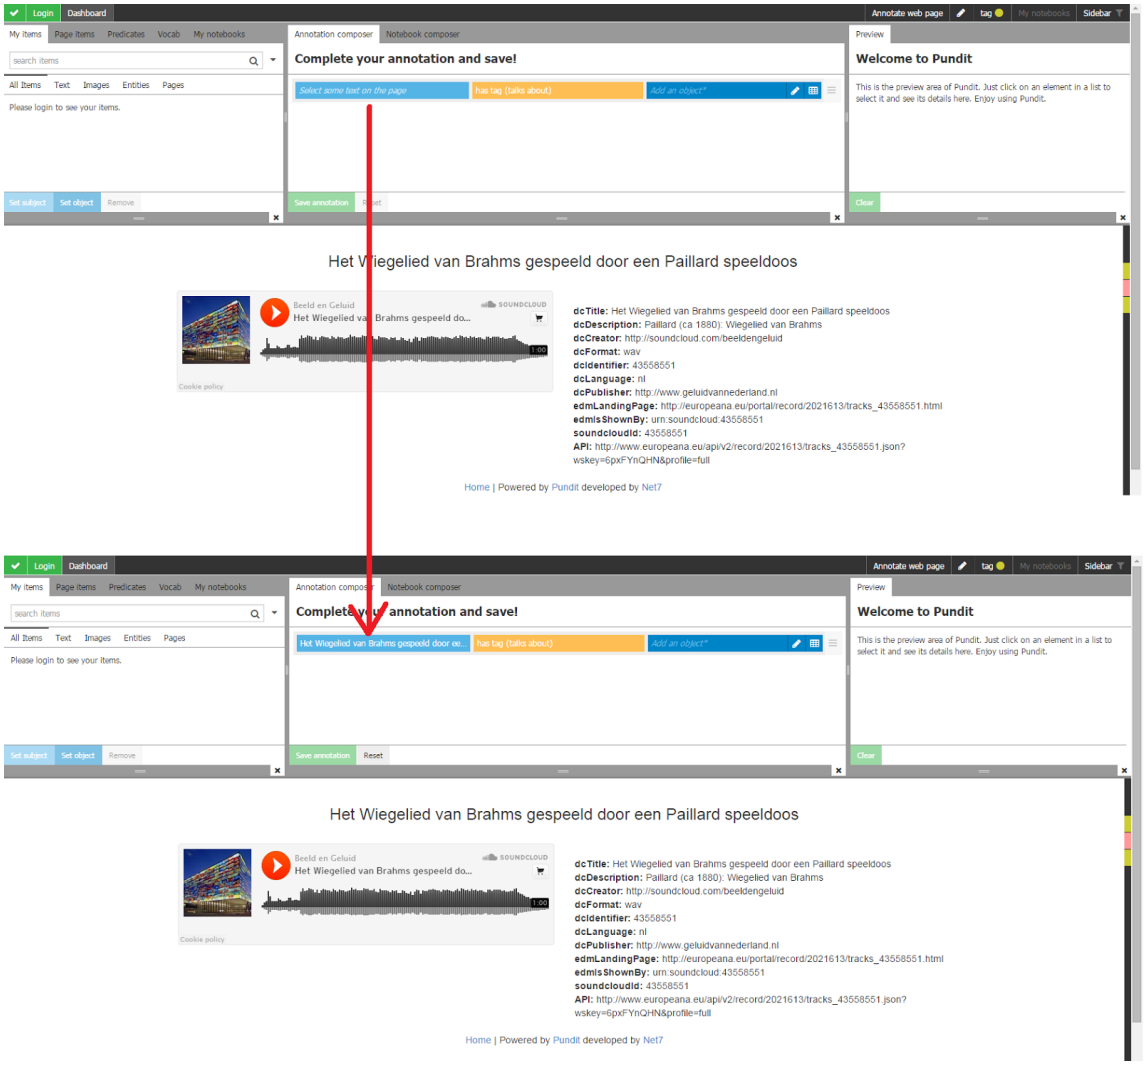

Image 32: Two screenshots of the widget, on top a screenshot with only the Dashboard open, underneath the screenshot after clicking 'Annotate web page'

The same confusion emerged when users' clicked 'Use' after choosing a concept (like topic 'Johannes Brahms') in the menu that appears when clicking 'Add an object*'. 'I don't see where it is going now', in order to properly create the annotation the user needs to also click 'Save annotation' after clicking 'Use' (See Image 33 underneath).

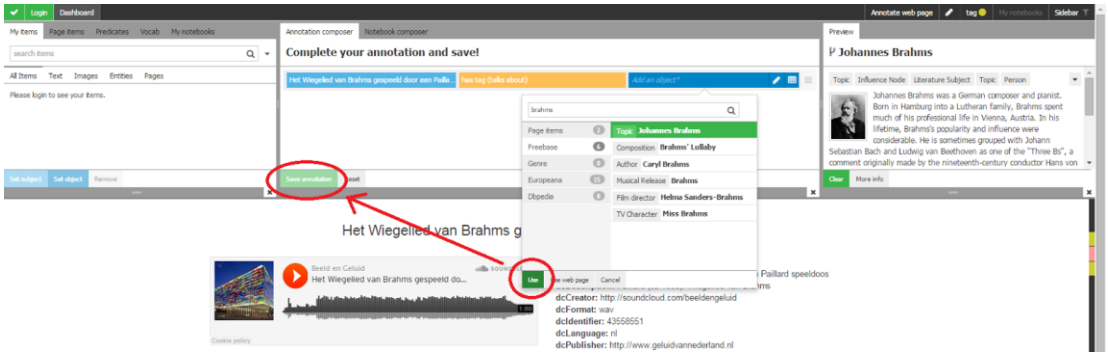

Image 33: Screenshot of widget when adding a concept from Freebase

An implicit result from the fact that users' could not always notice when an action was performed properly is that they also did not understand when they were clicking non-clickable options (because feedback in that case - of course - also lacked). For instance, one user did not understand that he did not actually save an annotation when clicking on 'Save annotation', because the 'Add an object' action was skipped. In this specific case the light green colour of the 'Save annotation' button (instead of dark green) was not enough to make clear this was not an option at that moment. One participant mentioned that there was no information about buttons and menus when hovering over the different elements in the interface.

Overall, the feedback of the different elements in the interface (buttons, search fields and menus) needs further refinement for a new version of the widget. A possible way to improve this is to create more contrast between the colours of the same button before and during hovering (and clicking). So instead of making buttons go from a light shade of green to a dark shade of green, change the colour completely if a button gets clicked on or when the cursor is hovering the button. The BBC for example uses more contrast in the colours of their buttons in the World Service Radio Archive (prototype) before and during hovering different elements, in some cases they also underline the clickable text when hovering. The image underneath shows the difference between Pundit and the World Service Radio Archive (prototype) in terms of hovering and not hovering over clickable elements and makes it clear how the visual feedback differs between the two platforms.

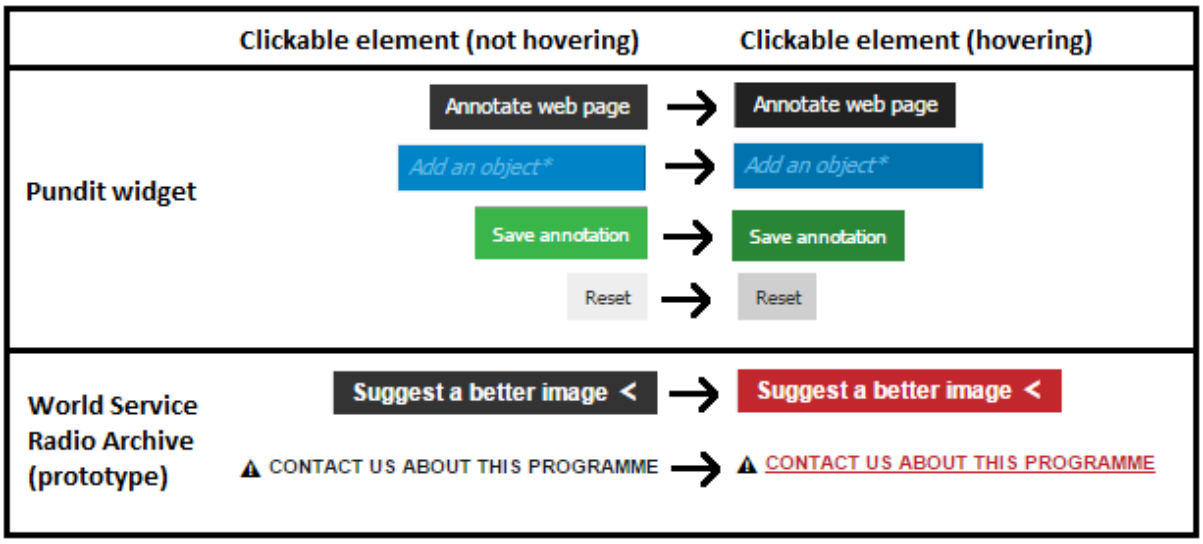

Image 34: Differences in visual feedback of various elements between the Pundit widget and the BBC's World Service Radio Archive (prototype).

3.1.5 Language and wording of interface elements

A lot of users had problems with understanding what the options of the several buttons represented. Part of the misunderstanding was caused by the name of the buttons. For some users it was not clear what is meant with "object" in the button "Add an object*", especially since they thought the object was the audio item in the SoundCloud player. One user agreed that it makes

sense to name it “object” when referring to an Europeana digital cultural object that is linkable, but suggested to use a word as ‘concept’ for this particular menu. The button “Has tag” also was problematic for users as well, ranging from questions if this should be a hashtag or if it means that a tag is already in place? The light blue box that says “Select some text on the page” made two users to try drag and drop pieces of text from the metadata next to SoundCloud player to the light blue box, which was not a supported action, and hence did not result in the behaviour intended by the user. In the task to delete a self-created annotation a couple of users were afraid of deleting too much information, because it was not clear what they would be deleted by clicking “Delete”. One user suggested to use “Remove”, because “Delete sounds so definitive, I think “Remove” would have been a better word, since it is just something I added myself”. Furthermore, two participants explicitly mentioned that they wanted the widget to be in Dutch (their native language).

To conclude, the user research showed that the choice of words and the used language in the widget needs further refinement. The metadata on Jocondelab is for example available in 14 different languages, using a bold font to indicate the displayed language in the top bar. Similar to the widget made by Pundit, Jocondelab[31] uses a fixed repository (Wikipedia) for crowdsourcing keywords. JocondeLab uses this repository to make available the items’ keywords in different languages.

Image 35: The 14 different languages in which you can view JocondeLab are listed in the top bar of the screen

3.1.6 Lack of guidance

Before the goal of the widget was explained to the participants in part 6, several questions regarding the widget were asked. Two participants guessed correctly that it is (also) possible to add something to the audio object with the widget bar in the form of a comment or a tag. Most participants called the widget a navigation bar to “see where you were the last time you visit the website” (participant 2), to ‘edit things or save things, sharing sounds and that it will help you search the database if you need sounds’ (participant 5), to “save, print, scrolling to next page” (participant 6) or to “share website with others” (participant 7).

[31] See: http://jocondelab.iri-research.org

Three of the participants explicitly mentioned that they were not invited actively by the widget: “I do not see a “help us describe” or other trigger to activate me to do something” (participant 1), “it isn’t really inviting to do something with this bar, I want to get an invite to really start working with this widget” (participant 7), “the widget can be better by triggering users more the first time they use it” (participant 4). The other participants will most likely share this opinion taken from the way they used the widget to try to complete five tasks in which they actually needed to add relevant information to the audio object. See the results of these five tasks in the figure underneath, task 6 was a free task to add whatever they could to the audio object, task 7 to add a comment, task 8 to add a genre, task 9 to link a relevant object from Europeana to the audio object and task 10 to use Freebase and/or DBpedia to enrich the audio object.

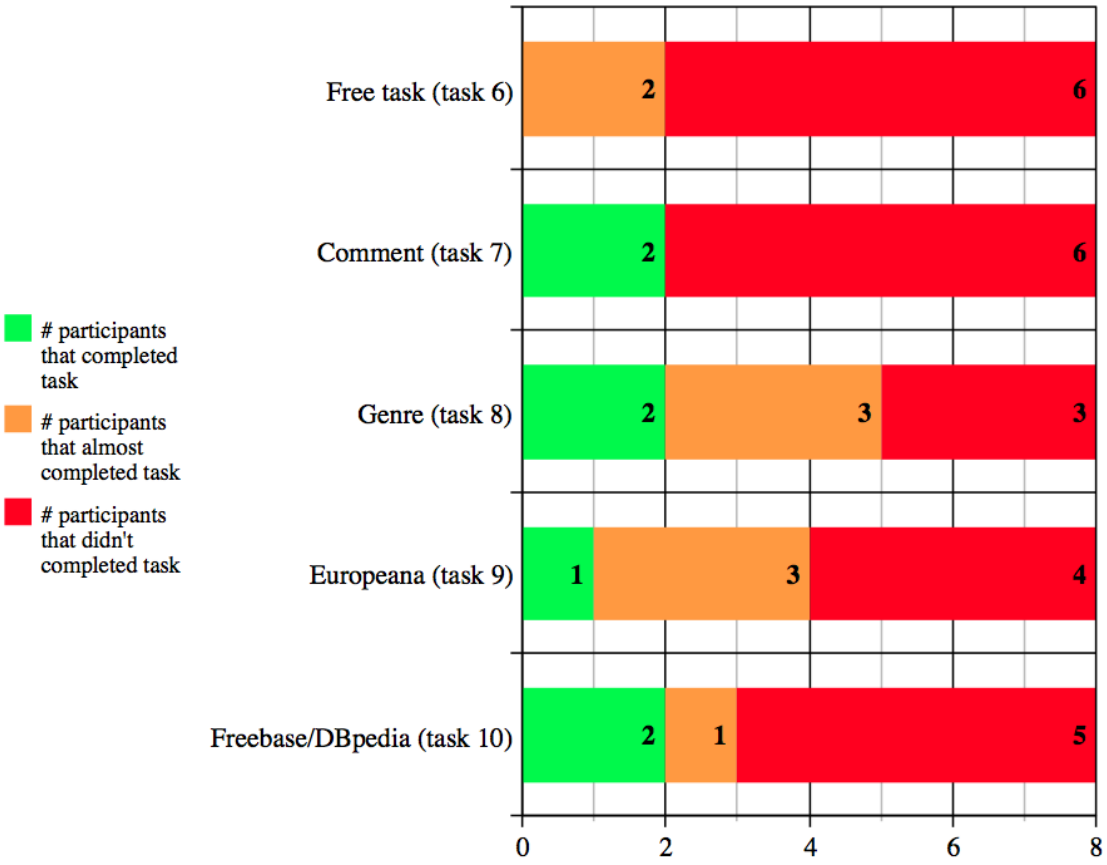

Image 36: Results of the level of completion of the enrichment tasks by all participants

The results shown in Image 36 indicate that most participants in all the tasks had not completed the task correctly (red). Sometimes participants came very close to completing the task (orange), but then got stuck just before they actually saved their annotation in the widget. There were three participants in total that executed at least one of the tasks correctly (green).

These results can be interpreted as an indicator to give more guidance to future users of the widget, the way to give this guidance could be to invite users actively to enrich audio objects within this widget as suggested by three of the participants. A possible way of facilitating this guidance and inviting users more actively is to embed a tutorial video to explain the functionalities of the website.

The crowdsourcing website NYPL Map Warper[32] (created for the New York Library) chooses to do so by embedding a tutorial video from YouTube on their Home-page.

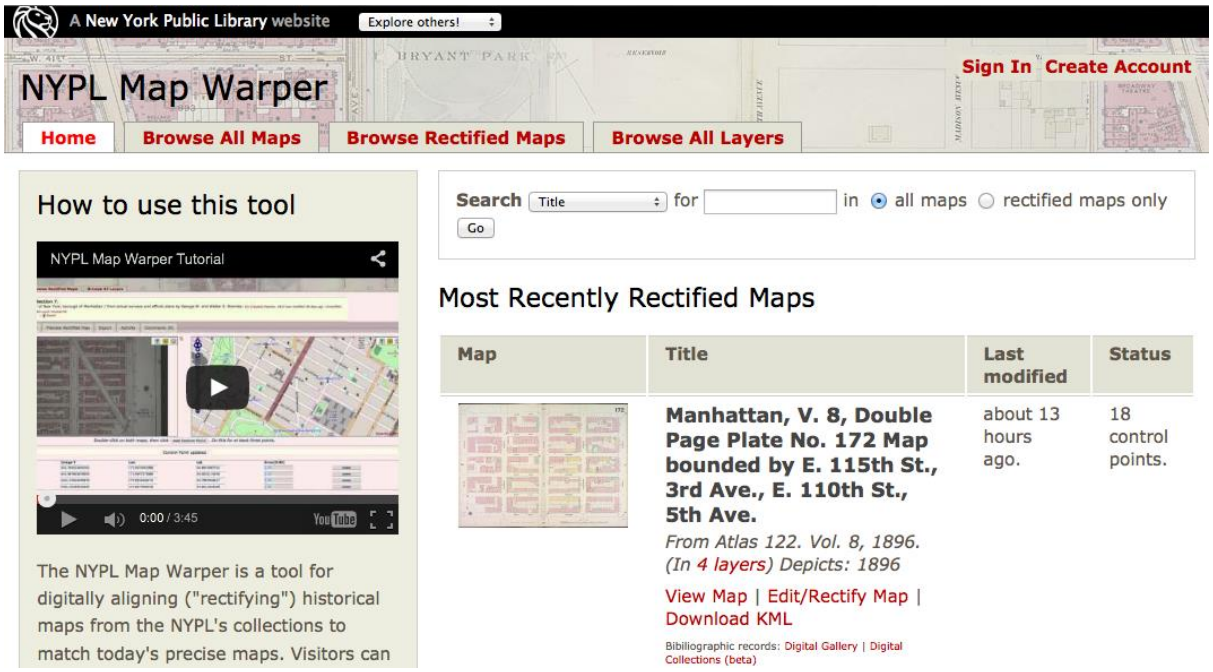

Image 37: NYPL Map Warper uses an embedded video tutorial to guide users through their crowdsourcing tool

Another possibility is to make the explanation of the widget available in text on the website, for instance behind a Help, About or Question Mark button. Next to these solutions which present the users with an extensive guide to the whole process, more subtle ways of providing guidance could be given as well. For example providing extra information when hovering over a button that has specific crowdsource functionality, like the example underneath on the JocondeLab website.

Image 38: Example of providing extra information when hovering over a button on the website of JocondeLab.

3.1.7 Lessons-learned and recommendations

From the four different problem areas that emerged during the user research, a couple of recommendations for further refinement of the widget are given:

[32] See: http://maps.nypl.org/

There was a general impression by the participants that the search field to search all the repositories stays hidden, expressed explicitly in words by some and implicitly by all participants when executing the different tasks. A recommendation to improve this problem area is (1) to consider what and where would be the first search field that users come across in the new version of the widget.

The problems concerning the feedback that was expected by the users after clicking buttons was addressed by most participants during the different tasks and taken together with the poor score of the tasks concerning the actual enriching of the audio object, provided an indicator (2) to improve the feedback of the buttons and menus.

Furthermore the wording of the interface elements, like buttons and fields names, was also a problem area for almost all users, besides two participants explicitly mentioned that they wanted the widget to be in their native language. Thus a recommendation is that (3) the word choices and the used language in the widget needs further refinement.

The last problem area emerged during the five tasks involving the actual enrichment of the audio object in which all participants experienced a lack of guidance from the widget when trying to enrich the audio object. The poor results of the completion of these five tasks can be interpreted as an indicator (4) to give more guidance in the new version of the widget, a possible way to give this guidance could be to invite users actively to enrich audio objects within this widget.

The sections above provide suggestions for possible solutions for the problem areas that emerged from the user evaluation. For continuously improving the widget over the course of the project , qualitative user research will be conducted by NISV after every major iteration. Giving specific tasks to test subjects that have affinity with sound related material, and identifying problem areas from these evaluation sessions will give important input to improve the widget fundamentally. The specific tasks will be created by NISV in consultation with NET7, in order to give focussed feedback next to the general outcome of the evaluation.

3.2 Pilot development

3.2.1 Pundit widget customized for Europeana Sounds

Through the Europeana search API, the metadata resources, identified by their Europeana ID, are retrieved. And a simple HTML page has been created to make it human readable. The application was built using the PHP language through the use of a parameter of type “GET” (Europeana Resource URI). The HTML page that is created includes and runs the Pundit client appropriately configured creating an annotation environment. For the pilot this annotation environment includes enrichment options based on Europeana, DBPedia, Freebase, and the Europeana Sounds SKOS Genre vocabulary (Ref 4). This application represents the working basis for the crowdsourcing widget that can be installed on the content provider of the Europeana Sounds project. The pilot widget can be found at the Following URL: http://euspndwidget.netseven.it

3.2.2 Lessons-learned and recommendations

To improve the overall user experience with Pundit, NET7 will improve the usability issues that have been highlighted during the user evaluation described in the previous chapter. The user testing reports have been analysed by NET7 to identify some recurring patterns concerning issues and problems the users encountered while completing their enrichment tasks.

Upcoming improvements to the Pundit widget will focus mainly on these aspects:

- Resource Panel: Pundit must provide more information to the user when there are no results and the panel is empty. And more importantly, create a version of the widget in which users immediately can type something in the resource panel after they indicated that they want to annotate the audio resource with a tag (in order to make it easier for users to search the different databases).
- General/visual feedbacks: The widget should always provide a visual effect after and during actions performed by the user (hovering and clicking), in order to provide clear feedback of what he/she is doing and how the system is responding. Furthermore, the widget also needs to give a good visual indication for all buttons and menus that are clickable and non-clickable, and make a clear visual distinction between the two.
- Saving annotations: Some users create annotations but forget to click the “Save annotation” button. A new version of the Pundit widget will provide a more evident call-to-action, or should save the annotation one step earlier, after the user clicks the ‘Use’ button. The name of this ‘Use’ button needs to change accordingly (for example: ‘Save tag’ and ‘Save comment’). The possibility to automatically save annotations while they are being created will also be considered.
- Web page annotation: The button that initiates the Pundit widget needs to be more context-specific (e.g. “Annotate the audio resource”).
- Notifications system: NET7 will introduce a notifications system to Pundit, in order to give some more clear and evident feedback about events to the user.
- All panels: Information must be provided explicitly about all the panels and their functionalities. This can be executed by creating a tutorial, an about or help section, which can be complemented with added information about the functionality in the form of a pop-up text when hovering over the panels.
- Preview panel: Users generally do not understand the function of the “Clear” button. We should change the micro-copy or re-think how the panel works.
- Page items: This panel will be removed from new versions of the Pundit widget, since it creates too much confusion for users.
- Templates: The “Pencil” icon is confusing users since it is usually related to the action of writing. An alternative icon needs to be introduced, based on a more suitable metaphor.

While working towards a version 1.0 release of the Pundit widget in January 2016, for use on their own websites by Data Providers, NET7 and NISV are planning more pilots to integrate more of the must-have crowdsourcing tasks, do more user evaluation and produce Guidelines on how to integrate the widget in external platforms maintained by the Data Providers.

To further develop the crowdsourcing widget and its application to sound-related cultural heritage objects, and to get as much Data Providers involved in the current development process and test phase WP2 utilized the WP2 session during the Training Days in Athens end of June 2015 to demonstrate the current work and its possibilities and then consult the Data Providers on their information needs and the possibilities with their collections (in terms of enriching them). One use-case that surfaced from this workshop was a shared interest, voiced by most Data Providers, in enriching their collections with musical instruments, preferably using a controlled vocabulary. This resulted in the plan to focus the next pilot with the Pundit widget on enriching multiple collections with concepts from the MIMO Vocabulary and Thesaurus[33]. This use-case seems to fulfil a real need from Data Providers and promises to provide a pilot with uptake by several Data Providers with diverse collections.

Image 39: Discussing the Pundit widget crowdsourcing features with the project partners in Athens

[33] See: http://www.mimo-international.com/vocabulary.html

4 Lessons-learned and recommendations

This section summarizes the lessons-learned and recommendation resulting from the work conducted in the context of the Crowdsourcing Infrastructure up until now.

4.1 System architecture and core infrastructure

When developing a set of crowdsourcing tools of which the majority will be integrated into one coherent infrastructure, it is vital to align all efforts in every stage of the project. This refers to things such as:

- Aligning all the work on data modelling and specifications in order to make sure data can be exchanged with the least technical efforts as possible (e.g. using a standard format for transferring annotations between Pundit and the Annotations API).
- Working from the start with an integration perspective and iterative development (releases) in mind: test the integration at every stage. Europeana can provide guidelines and test environments to work against.

Part of the work under the Taskforce on Enrichment and Evaluation Methods, organized in the context of EuropeanaTech (and with WP2 being represented by WP lead Maarten Brinkerink, as announced in D2.2), has been the creation of an annotated corpus of metadata records (also known as 'Gold Standard', Ref 3), against which the results of an automatic enrichment tool can be evaluated to determine its qualitative performance. While creating this corpus, the task force recognized that the task of annotating the corpus, and also determining the correctness of the annotations, could be a relevant application and testing scenario for the Europeana Sounds Crowdsourcing Infrastructure. Looking at the Epics described in D2.2 (Ref 1), it makes sense to enhance the *Object tagging* epic (E2) with a specialization to support semantic enrichments (following the proposal from the PATH project to model semantic enrichments as annotations using the Open Annotation model, which has been adopted by WP2), and *Moderation epic* (E9) for reaching the agreement and confirmation that the annotation is deemed correct.

4.2 Traditional music pilot

Shift is aiming to help traditional musicians engage with archive collections containing traditional music more frequently and more deeply, by connecting websites popular with traditional musicians to the collections of multiple archives. By linking archives and traditional musicians using digital tools, we can reduce barriers in teaching and learning. This also maximises the possibility for reuse of the metadata aggregated within the Europeana Sounds project, and makes discovery and reuse of archival materials more meaningful for specific knowledge communities.

By engaging with the active traditional music community through surveys and social media, Shift has been able to define a particular need, i.e. a quicker way of identifying tunes and a more intuitive way of exploring traditional music archives.

As seen in Athens, other Data Providers who do not have traditional Irish and Scottish archival material were interested to see how they could use the learning developed as part of the traditional music pilot. Extrapolating the learning from this pilot into general lessons learned about archives and their audiences, and sharing this with the Europeana Sounds project, will be a focus point for Shift in the second half of the project.

By working together with the developer of an existing tool that already meets some of these needs, Shift is hoping to create a lasting legacy for the Europeana Sounds project, in the shape of an open-source app that uses the Europeana API to help users explore sound archives.

4.3 Crowdsourcing widget

From the four different problem areas that emerged during the user research, a couple of recommendations for further refinement of the widget are given:

There was a general impression by the participants that the search field to search all the repositories stays hidden, expressed explicitly in words by some and implicitly by all participants when executing the different tasks. A recommendation to improve this problem area is (1) to consider what and where would be the first search field that users come across in the new version of the widget.

The problems concerning the feedback that was expected by the users after clicking buttons was addressed by most participants during the different tasks and taken together with the poor score of the tasks concerning the actual enriching of the audio object, provided an indicator (2) to improve the feedback of the buttons and menus.

Furthermore the wording of the interface elements, like buttons and fields names, was also a problem area for almost all users, besides two participants explicitly mentioned that they wanted the widget to be in their native language. Thus a recommendation is that (3) the word choices and the used language in the widget needs further refinement.

The last problem area emerged during the five tasks involving the actual enrichment of the audio object in which all participants experienced a lack of guidance from the widget when trying to enrich the audio object. The poor results of the completion of these five tasks can be interpreted as an indicator (4) to give more guidance in the new version of the widget, a possible way to give this guidance could be to invite users actively to enrich audio objects within this widget.

To improve the overall user experience with Pundit, NET7 will improve the usability issues that have been highlighted during the user evaluation. The user testing reports have been analysed by NET7 to identify some recurring patterns concerning issues and problems the users encountered while completing their enrichment tasks.

For continuously improving the widget over the course of the project qualitative user research will be conducted by NISV after every major iteration. Giving specific tasks to test subjects that have affinity with sound related material, and identifying problem areas from these evaluation sessions will give important input to improve the widget fundamentally. The specific tasks will be created by

NISV in consultation with NET7, in order to give focussed feedback next to the general outcome of the evaluation.

References

Ref 1	D2.2: Functional design of semantic enrichment http://pro.europeana.eu/files/Europeana_Professional/Projects/Project_list/Europeana_Sounds/Deliverables/EuropeanaSounds-D2.2-Functional-design-of-semantic-enrichment-v1.0.pdf
Ref 2	MS11: Evaluation of first deployment of the crowdsourcing infrastructure https://basecamp.com/1936492/projects/2141070/messages/43206834?enlarge=158993950#attachment_158993950
Ref 3	On How to Perform a Gold Standard Based Evaluation of Ontology Learning Klaas Dellschaft and Steffen Staab http://iswc2006.semanticweb.org/items/Dellschaft2006fy.pdf
Ref 4	D1.3: Ontologies for Sound http://pro.europeana.eu/files/Europeana_Professional/Projects/Project_list/Europeana_Sounds/Deliverables/EuropeanaSounds-D1.3-Ontologies-for-sound-v1.2.pdf
Ref 5	EUscreenXL the pan Europeana audiovisual aggregator for Europeana: D4.4. Pilot development
Ref 6	Europeana Space, Spaces of possibility for the creative re-use of digital cultural content: D2.1 Requirements for the creative use of Digital Cultural Resources; progress on collaboration towards Europeana Labs

Appendix I: User research for pilot application as part of Europeana Sounds

Respondents: 59

Running time: March 2015

As part of our work on Europeana Sounds, Historypin is creating a TuneID-o-Matic, a website-based app into which you can play a traditional Irish or Scottish tune and it surfaces archival recordings of that tune, pulled from the Europeana database. These tunes are supplied to Europeana by Tobar an Dualchais and Comhaltas.

We think that such an app will lead to more discovery of archival material related to Irish and Scottish traditional music and to greater enjoyment among traditional musicians as they can find authoritative sources for tune names and other metadata, as well as different versions of the same tune, in an easier and more engaging way.

To inform the technology brief, we needed to do some research into the needs and wants of traditional musicians, our target audience.

Based on anecdotal evidence from conversations with musicians and discussions on the traditional music forum www.thesession.org, it seems that the main question that traditional musicians are trying to answer is: What is the name of this tune? Knowing a tune's name will enable you to identify versions of it in the corpus of traditional music, as well as to find out its history. It will also give you access to either recordings or transcriptions, which can help you learn the tune.

To understand where traditional musicians hear new tunes, how they identify these and what resources and tools they subsequently use to learn the tune, Historypin created a survey, which was posted online and distributed among traditional musicians on Twitter and through word of mouth. In addition, Historypin also interviewed two traditional musicians by phone and entered their comments into the survey form.

The survey was active throughout March 2015 and attracted 59 responses in total. The traditional musicians who responded came from a variety of places, with the most coming from the USA and Canada. Many of them play a range of instruments, and most of them attend local sessions.

Where do people mostly hear new tunes?

1. Other players (either at a session or a feis)
2. CDs
3. YouTube
4. Radio / podcasts

Workflows for learning unfamiliar tunes:

1. RECORDING

Record a tune at a session or concert → write it out in ABC or staff notation → look it up in tune books or online → learn tune from notation

Record a tune at a session or concert → look it up online (using Tunepal?) → find name, notation, recordings etc. → learn tune from any of these sources

2. FIND NAME (below two workflows are most common among survey respondents)

Ask name of tune → look up recordings of the tune → learn tune from the recording

Ask name of tune → look up notation for the tune → learn tune from notation

3. LEARN BY EAR

Listen to a tune at a session → try to follow along and learn it by ear

Listen to a tune from a recording → try to follow along and learn it by ear

4. ASK A FELLOW MUSICIAN

Ask a fellow musician to teach the tune to you → learn it by ear

Sources people use for finding tunes:

- their own collection, either on the computer or on cds
- online sources dedicated to trad like TheSession, JCs Tunebook, the Reverend Tune Finder, Henrik Norbeck ABC tunes, EasyABC, abctunesearch, Tunepal, Chiff and Fipple, irishtune.info, John Chambers Tune Finder
- physical tunebooks like the Breatnach Books 1-5, Bulmer Sharpley 1-4, O'Neills, Goodman, Petrie, Comhaltas session tunebooks
- online sources not specifically dedicated to trad like YouTube, SoundCloud, Spotify
- radio stations like Clare FM

Recording tunes:

If musicians record tunes, they usually do this either on a handheld recording device or on a smart phone.

Transcribing tunes:

Musicians either transcribe in staff notation or in ABC. Many use online transcription tools to help them with this.

Tools:

Many musicians use the Amazing Slow Downer App to slow tune recordings down so it's easier to learn them by ear.

User persona

Actual response:

“long process....at a session, i record it so that i can learn the way it's played at the session. I might Tunepal it later at home to find out a name, and I search for recordings of it and listen to them tirelessly in the car & home for a few weeks. I use thesession.org to find what recordings are out there & search for them--often picking recordings of players I know or like. Then I pick a version I like (or a combination), and if it works with what is played at session, I learn that. If the session version is any different, I probably learn them both and challenge myself to remember to play it differently at session.”

This musician seems to fit our ideas best as a user persona for the TuneID-o-Matic

They:

1. record the tune to learn it
2. upload it into Tunepal it to find a name
3. then use the name to search for recordings of that tune
4. then pick a recording they like best and learn that

We propose to combine step 2 and 3 into the Tune-ID-o-Matic functionality using the archival recordings that are uploaded to Europeana as part of Sounds.

Do you attend any sessions?

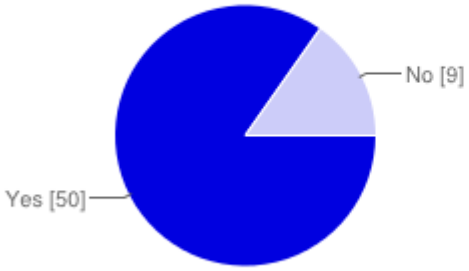

Yes	**50**	84.7%
No	**9**	15.3%

Do you record tunes?

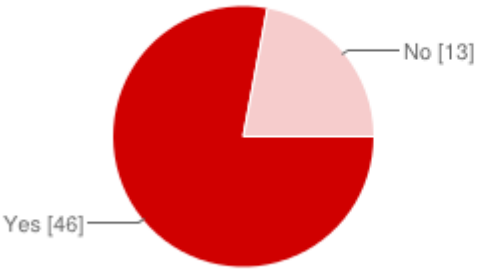

Yes	**46**	78%
No	**13**	22%

Do you transcribe tunes?

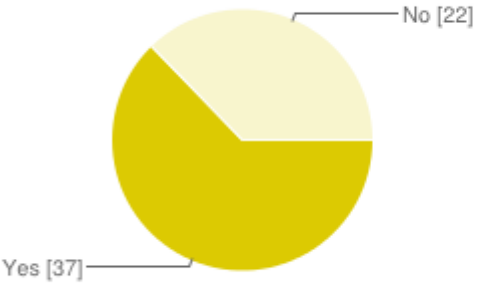

Yes	**37**	62.7%
No	**22**	37.3%

Questionnaire:

1. What instruments do you play?
2. How did you learn to play traditional music?
3. Do you attend any sessions?
4. Could you tell us a bit about your session?
5. Could you describe what generally happens during a typical session?
6. What do you like best about your session?
7. Would it be okay if we came and observed your session sometime?
8. Where do you hear new tunes?

9. What do you do when you hear an unfamiliar tune that you'd like to learn?
10. What sources do you use for this?
11. Do you record tunes?
12. How do you do that?
13. What do you do with the recordings?
14. Do you transcribe tunes?
15. How do you do that?
16. What are the reasons for transcribing tunes?
17. Do you know any traditional music archives?
18. Which traditional music archives do you know?
19. Do you use them, and if so, how do you use them?
20. Do you visit any websites related to traditional music?
21. Which ones?
22. What do you normally use them for?
23. Is there anything else you'd like to share with us about playing and learning traditional music?
24. Would you be happy for us to get in touch with you if we have any further questions?

Appendix II: Questions for the tasks for the crowdsourcing widget

Underneath the tasks and questions that were assigned and asked by the Netherlands Institute for Sound and Vision for the user research of the widget created by Net7. Originally all in Dutch, but provided with an English translation as well for this Milestone.

1. Opdracht: Login (2 minuten)

1. Task: Login (2 minutes)

Opdracht: Ga naar http://euspndwidget.netseven.it/index.php?id=2021613%2Ftracks_43558551 en log in.

Task: Go to http://euspndwidget.netseven.it/index.php?id=2021613%2Ftracks_43558551 and try to login.

- *Wat doet de deelnemer?*

- *What is the participant doing?*

2. Opdracht: Ontdek website (4 minuten)

2. Task: Explore website (4 minutes)

Opdracht: Alles onder de zwarte bovenste balk moet de website van Beeld en Geluid voorstellen waar je een item uit onze collectie ziet. Nog nergens op drukken en hardop vertellen wat je ziet.

Task: Everything underneath the top dark bar is meant to visualise the website of the Netherlands Institute of Sound and Vision where you see one of the items of our collection. Don't click on anything and tell us what you see.

(als deelnemer het heeft over knoppen op website, vragen wat ze denken dat er gebeurd als ze er op zouden klikken).

(if the participant is mentioning buttons on the website, ask what they think will happen if they should click on it).

- *Paginastructuur (structuur van de onderdelen op de pagina)*

- *Page structure (structure of the different elements on the page)*

- *Gebruiksvriendelijkheid / Interface*

- *Userfriendliness / Interface*

- *Look / Feel*

- *Look / Feel*

3. Vraag: Wat denk je hier te kunnen doen? (2 minuten)

3. Question: What do you think you can do here? (2 minutes)

4. Vraag: Wat zou het eerste zijn waar je op zou drukken? (3 minuten)

4.Question: What would be the first thing you would to click on? (3 minutes)

Waarom?

Why?

Wat denk je dat er zal gebeuren?

What do you think that shall happen?

- *Paginastructuur (structuur van de onderdelen op de pagina)*

- *Page structure (structure of the different elements on the page)*

- *Gebruiksvriendelijkheid / Interface*

- *Userfriendliness / Interface*

- *Look / Feel*

- *Look / Feel*

5. Opdracht: Klikken (3 minuten)

5. Task: Click (3 minutes)

Ok, je mag nu klikken waar je op zou klikken, vertel hardop wat je verwacht en wat er vervolgens gebeurt.

Ok, you can now click, tell us what you are expecting and what you think actually happens after clicking.

- *Wat doet de deelnemer?*

- *What is the participant doing?*

- *Paginastructuur (structuur van de onderdelen op de pagina)*

- *Page structure (structure of the different elements on the page)*

- *Gebruiksvriendelijkheid / Interface*

- *Userfriendliness / Interface*

- *Look / Feel*

- Look / Feel

6. Vraag: Doel (8 minuten)

6. Question: Goal (8 minutes)

De widget is gemaakt om digitale objecten van Beeld en Geluid te annoteren (verrijken) en om dergelijke objecten te linken met relevante digitale objecten van Europeana, dit om de collectie van Beeld en Geluid gekoppeld wordt aan trefwoorden die internationaal gebruikt worden, wat onze collectie beter vindbaar en makkelijker doorzoekbaar maakt. Hoe zou je dit proberen te doen zonder onze hulp?

The widget is created to annotate digital objects of Sound and Vision (enriching) and to link these objects with relevant other digital objects in the collection of Europeana, this is made to link the collection of Sound and Vision to internationally used key words, that makes our collection better searchable and makes it easier to find for other people. Now you know the goal of this widget, how should you try to do this without our help?

- Wat doet de deelnemer?

- What is the participant doing?

- Paginastructuur (structuur van de onderdelen op de pagina)

- Page structure (structure of the different elements on the page)

- Gebruiksvriendelijkheid / Interface

- Userfriendliness / Interface

- Look / Feel

- Look / Feel

7. Opdracht: Voeg een (persoonlijke) opmerking toe aan het audio object (7 minuten)

7. Task: Add a (personal) comment to the audio object (7 minutes)

Opdracht: Luister naar het audio fragment en voeg een opmerking toe.

Task: Listen to the audio object and add a comment.

- Wat doet de deelnemer?

- What is the participant doing?

- Paginastructuur (structuur van de onderdelen op de pagina)

- Page structure (structure of the different elements on the page)

- Gebruiksvriendelijkheid / Interface

- Userfriendliness / Interface

- Look / Feel

- Look / Feel

8. Opdracht: Genre toevoegen (5 minuten)

8. Task: Add genre (5 minutes)

Opdracht: Voeg een relevant genre toe aan het audio fragment

Task: Add a relevant genre to the audio object

- Wat doet de deelnemer?

- What is the participant doing?

(zodra ze in het tag-menu zijn: er achter komen of ze snappen dat de preview window slaat op het uitklapmenu) > JA/NEE?

(as soon participants are in the tag menu: check if they understand that the preview window is linked to the concepts in the menu) > YES/NO?

- Paginastructuur (structuur van de onderdelen op de pagina)

- Page structure (structure of the different elements on the page)

- Gebruiksvriendelijkheid / Interface

- Userfriendliness / Interface

- Look / Feel

- Look / Feel

9. Opdracht: Linken (5 minuten)

9. Task: Linking (5 minutes)

Opdracht: Link een relevant digitaal object van Europeana aan het audio fragment.

Task: Link a relevant digital object from the collection of Europeana to the audio object.

- Wat doet de deelnemer?

- What is the participant doing?

- Paginastructuur (structuur van de onderdelen op de pagina)

- Page structure (structure of the different elements on the page)

- Gebruiksvriendelijkheid / Interface

- Userfriendliness / Interface

- Look / Feel

- Look / Feel

10. Opdracht: DBpedia/Freebase (5 minuten)

10. Task: DBpedia/Freebase (5 minutes)

Opdracht: Voeg relevante informatie toe aan het audio object door gebruik te maken van DBpedia en Freebase.

Task: Add relevant information to the audio object by using DBpedia and Freebase

- Wat doet de deelnemer?

- What is the participant doing?

Bekijk vervolgens ook de originele informatie op Freebase of DBpedia van het concept dat je hebt toegevoegd aan het audio fragment.

Now also search for the original information on the website of Freebase or DBpedia of the concept that you added to the audio object.

- Wat doet de deelnemer?

- What is the participant doing?

- Paginastructuur (structuur van de onderdelen op de pagina)

- Page structure (structure of the different elements on the page)

- Gebruiksvriendelijkheid / Interface

- Userfriendliness / Interface

- Look / Feel

- Look / Feel

11. Opdracht: Verwijder een eerder gemaakte toevoeging (3 minuten)

11. Task: Delete an annotation (3 minutes)

- Wat doet de deelnemer?

- What is the participant doing?

- Paginastructuur (structuur van de onderdelen op de pagina)

- Page structure (structure of the different elements on the page)

- Gebruiksvriendelijkheid / Interface

- Userfriendliness / Interface

- Look / Feel

- Look / Feel

12. Conclusie (8 minuten)

12. Conclusion (8 minutes)

Geef op onderstaande onderdelen een cijfer van 1 (heel slecht) tot 10 (heel goed) en noem minimaal een reden waarom je dit cijfer hebt gegeven.

Give the following elements a grade between 1 (very bad) till 10 (very good) and name at least one reason why you've given this grade.

- Is de widget gebruiksvriendelijk?
- Is the widget userfriendly?
- Waren de knoppen makkelijk te vinden? (paginastructuur)
- How easy was it to find the buttons? (page structure)
- Het concept van de widget, wat je er uiteindelijk mee kan doen? (functionaliteiten)
- Concepts of the widget, what you can do with it eventually? (functionalities)

Overige vragen:

Other questions:

- Wat miste je qua dingen die er mogelijk waren als je bedenkt dat dit een widget is die je helpt een digitaal cultureel object te verrijken met jouw kennis? (functionaliteit)
- What did you miss in the widget when you know this is a tool that helps you to add your own knowledge to digital cultural objects? (functionalities)
- Wat waren de goede punten van deze widget?
- What were the good points of this widget?
- Wat kan er verbeterd worden aan de widget volgens jou?
- How can this widget be improved?
- Overige suggesties?
- Other suggestions?
- Nog op of aanmerkingen over het onderzoek zelf?
- Suggestions regarding the user research itself?
- En over de manier waarop het onderzoek is afgenomen?
- And about the way we did this user research?

Appendix III: Terminology

A project glossary is provided at: http://pro.europeana.eu/web/guest/glossary.

Additional terms are defined below:

Term	Definition
AB	Advisory Board
APEX	Archives Portal Europe network of excellence
EC-GA	Grant Agreement (including Annex I, the Description of Work) signed with the European Commission
GA	General Assembly
TEL	The European Library
UAP	User Advisory Panel
WP	Work Package