Polifonia: a digital harmoniser for musical heritage knowledge, H2020

D2.1: Ontology-based knowledge graphs for music objects (V1.0)

Deliverable information	
WP	WP2 Musical Heritage Knowledge Graphs
Deliverable dissemination level	PU Public
Deliverable type	R Document, report
Lead beneficiary	KCL
Contributors	KCL, UNIBO, NISV, IReMus, DANS
Document status	Final
Document version	V1.0
Date	20/12/2021
Authors	Albert Meroño-Peñuela (KCL), Jacopo de Berardinis (KCL), Delfina Sol Martinez Pandiani (UNIBO), Valentina Anita Carriero (UNIBO), Mari Wigham (NISV), Andrea Poltronieri (UNIBO), Fiorela Ciroku (UNIBO), Christophe Guillotel-Nothmann (IReMus), Philippe Rigaux (IReMuS)

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101004746

PAGE INTENTIONALLY BLANK

Project Information

Project Start Date: 1st January 2021
Project Duration: 40 months
Project Website: https://polifonia-project.eu

Project Contacts

Project Coordinator

Valentina Presutti

ALMA MATER STUDIORUM - UNIVERSITÀ DI BOLOGNA
Department of Language, Literature and Modern Cultures (LILEC)

E-mail: valentina.presutti@unibo.it

Project Manager

Marta Clementi

ALMA MATER STUDIORUM - UNIVERSITÀ DI BOLOGNA
Research division

E-mail: marta.clementi3@unibo.it

POLIFONIA Consortium

No.	Short name	Institution name	Country
1	UNIBO	ALMA MATER STUDIORUM - UNIVERSITÀ DI BOLOGNA	Italy
2	OU	THE OPEN UNIVERSITY	United Kingdom
3	KCL	KING'S COLLEGE LONDON	United Kingdom
4	NUI GALWAY	NATIONAL UNIVERSITY OF IRELAND GALWAY	Ireland
5	MiC	MINISTERO DELLA CULTURA	Italy
6	CNRS	CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS	France
	SORBONNE	SORBONNE UNIVERSITE (LinkedTP)	France
7	CNAM	CONSERVATOIRE NATIONAL DES ARTS ET METIERS	France
8	NISV	STICHTING NEDERLANDS INSTITUUT VOOR BEELD EN GELUID	Netherlands
9	KNAW	KONINKLIJKE NEDERLANDSE AKADEMIE VAN WETENSCHAPPEN	Netherlands
10	DP	DIGITAL PATHS	Italy

Project Summary

European musical heritage is a dynamic historical flow of experiences, leaving heterogeneous traces that are difficult to capture, connect, access, interpret, and valorise. Computing technologies have the potential to shed a light on this wealth of resources by extracting, materialising and linking new knowledge from heterogeneous sources, hence revealing facts and experiences from hidden voices of the past. Polifonia makes this happen by building novel ways of inspecting, representing, and interacting with digital content. Memory institutions, scholars, and citizens will be able to navigate, explore, and discover multiple perspectives and stories about European Musical Heritage.

Polifonia focuses on European Musical Heritage, intended as musical contents and artefacts - or music objects - (tunes, scores, melodies, notations, etc.) along with relevant knowledge about them such as: their links to tangible objects (theatres, conservatoires, churches, etc.), their cultural and historical contexts, opinions and stories told by people having diverse social and artistic roles (scholars, writers, students, intellectuals, musicians, politicians, journalists, etc), and facts expressed in different styles and disciplines (memoire, reportage, news, biographies, reviews), different languages (English, Italian, French, Spanish, and German), and across centuries.

The overall goal of the project is to realise an ecosystem of computational methods and tools supporting discovery, extraction, encoding, interlinking, classification, exploration of, and access to, musical heritage knowledge on the Web. An equally important objective is to demonstrate that these tools improve the state of the art of Social Science and Humanities (SSH) methodologies. Hence their development is guided by, and continuously intertwined with, experiments and validations performed in real-world settings, identified by musical heritage stakeholders (both belonging to the Consortium and external supporters) such as cultural institutes and collection owners, historians of music, anthropologists and ethnomusicologists, linguists, etc.

Executive Summary

This deliverable describes the results of the WP2 Musical Heritage Knowledge Graphs in its first 12 months with respect to the task *T2.1: Ontology-based knowledge graphs for music objects*. The project, and specifically WP2 and T2.1, uses knowledge engineering as a process and methodology to build ontologies and knowledge graphs that facilitate large-scale integration of various sources of musical heritage data. In this deliverable, we focus on *music objects* broadly understood as artefacts carrying musical content and meaning (recordings, music notation, musical theory, musical features) rather than contextual descriptions (e.g. composer, year of performance). New conceptualisations, in the form of ontologies, are needed to formalise the semantics of this domain and allow the creation of interoperable knowledge graphs from datasets.

To achieve this, we apply a well-known ontology design methodology, eXtreme Design (XD), which emphasises: (i) ontology reuse and modularity through the design and use of Ontology Design Patterns (ODPs); and (ii) gathering a comprehensive inventory of ontological requirements through **competency questions** (CQs). To implement (i), we survey existing ontologies and ODP that are relevant to the domain, and existing datasets from Polifonia's WP1 Pilots – building on the dataset survey work from WP1 [1]– and from the Music Information Retrieval (MIR) community; to implement (ii) we follow a thorough process of collecting and normalising CQs from the stories and personas that stem from the Pilots of WP1. This ensures an agnostic approach in which Polifonia designs a specific conceptualisation of music (elements, reception, actors, events, etc.) that reflects original and innovative theory building aspirations, while providing a good compromise with respect to reusing existing models and theories. As a result of this, we report on the first version of the **Polifonia Ontology Network** (PON), a network of ontologies that model the domains emanating from such stories and personas in the *Core*, *Musical Performance*, *Musical Composition*, *Musical Feature*, *Source*, *Instrument*, *Comparative Measure*, *Music Emotion*, *Bell*, and *Metadata* ontology modules.

Even though the evaluation of this knowledge engineering work is the specific objective of another task (T2.5 and *D2.7, D2.8: Ontology testing and evaluation report*), we report here on preliminary efforts to evaluate the adequacy and correctness of the PON at effectively answering the project's CQs, through testing automation and expert validation; and its applicability in a real-world, demanding setting in which we used the developed ontologies to annotate, model and exchange data in a demonstration at a top, internationally renowned event: the **AI & Music Festival at the Sónar Barcelona International Music Festival** (27-28 October 2021, Barcelona).

Finally, we also pave the way into the project-wide, large-scale usage of the PON to annotate all the identified datasets in Polifonia (through [1] and here), MIR and beyond, effectively creating the **Polifonia Knowledge Graph** (PKG). The specific details on how this is achieved will follow in a further deliverable (*D2.4: Methods for interlinking knowledge graphs*), but we offer here a roadmap and the first methods we have created to interlink musical heritage knowledge graphs.

Document History

Version	Release date	Summary of changes	Author(s) - Institution
V0.1	16/09/2021	First draft released	Albert Meroño-Peñuela (KCL)
V0.2	28/10/2021	Second draft, task distribution	Albert Meroño-Peñuela (KCL), Jacopo de Berardinis (KCL), Delfina Sol Martinez Pandiani (UNIBO), Mari Wigham (NISV)
V0.3	04/11/2021	Third draft, almost complete	Albert Meroño-Peñuela (KCL), Jacopo de Berardinis (KCL), Valentina Anita Carriero (UNIBO), Andrea Poltronieri (UNIBO), Delfina Sol Martinez Pandiani (UNIBO), Mari Wigham (NISV)
V0.4	11/11/2021	Fourth draft	Jacopo de Berardinis (KCL), Valentina Anita Carriero (UNIBO), Andrea Poltronieri (UNIBO), Fiorela Ciroku (UNIBO), Christophe Guillotel-Nothmann (IReMus), Philippe Rigaux (IReMus), Andrea Scharnhorst (DANS)
V0.5	12/11/2021	Release for first round of review	Albert Meroño-Peñuela (KCL)
V0.6	13/12/2021	Fixes after first round of review	Albert Meroño-Peñuela (KCL), Jacopo de Berardinis (KCL), Valentina Carriero (UNIBO), Fiorela Ciroku (UNIBO), Andrea Poltronieri (UNIBO)
V1.0	20/12/2021	Final version submitted to EU	UNIBO

Table of contents

1 Introduction

This deliverable describes the first version of the ontologies and knowledge graphs of music objects and patterns in Polifonia; as well as the software engineered to this end. It contains a description of the ontologies developed for musical objects, patterns, and musical heritage metadata; and a preliminary technical report on a software package that transforms music collections –containing those aforementioned objects, patterns, and metadata– into knowledge graphs that use the developed ontologies as basic semantic schemas to organise their content. The specific point of departure of this deliverable are the existing ontologies and ontology design patterns (ODP) that are relevant to the domain of European musical cultural heritage, and the existing related datasets that were surveyed in, and requirements that stemmed from, previous Polifonia work in WP1 Pilots, and their related stories, personas and datasets [1]. These pilots, stories and personas constitute our modelling domain and main source of requirements.[1]

It is broadly acknowledged that key datasets for the understanding of European musical heritage are scattered on the Web, and are hard to find, reuse, integrate and combine to extract insights. These limitations hinder the reusability of datasets in European musical heritage, making them poorly suited for modern methodologies and frameworks supporting data reuse, for example FAIR [2]. One way of addressing these limitations is through the use of the so-called Semantic Web technology stack: a Web data publishing paradigm that relies on Linked Data [3] and ontologies [4] to generate Knowledge Graphs, large data representations that "use a graph-based data model to capture knowledge in application scenarios that involve integrating, managing and extracting value from diverse sources of data at large scale" [5]. These ontologies and Knowledge Graphs have been deployed in a number of specific areas of musical data with some success in their integration, albeit still in a reduced number compared to the overall availability of these datasets [6]. For example, the Music Ontology [7] and the DOREMUS ontology [8] describe high-level entities in the musical heritage domain, such as composers, bands, and performances; other ontologies, such as the Music Score Ontology [9] and the MIDI ontology [10] focus on the representation of the low-level, fine-grained entities present in symbolic music representations, such as notes and instruments. Other ontologies provide classes and properties that can be used for enriching the semantic annotations of an increasing number of specialised datasets. However, this growing body of work on musical ontologies is highly sparse, with many ontologies not linking or reusing classes or properties from any other; this contributes to increasing the sparseness of the conceptual space for interlinking music datasets. Consequently, there are two significant gaps in the current body of musical ontologies and the Polifonia objectives:

1. It is hard to decide what ontological resources can be used to model **the music object**

[1] An additional persona, Sonia, which was added to those of the pilots for the specific purposes of the Sónar demo (see Chapter 4), whose competency questions overlap approximately at an 80% with personas from the pilots.

itself[2], and which ones can be used to model **its cultural and historical context**; as well as which of those resources have been modelled in previous ontologies, and which ones need to be engineered from scratch

2. It is hard to assess the differences in the methodologies, purposes and competency questions that were used in the design of those ontologies; therefore complicating their further integration and compatibility

These gaps are exacerbated by the intrinsic challenge that comes with the **formalisation of musical theories that hold valid with independence from cultural and historical contexts, and the pass of time**. Indeed, this is a grand enterprise that is not novel: through the centuries, various acclaimed music theorists such as Gioseffo Zarlino (1517 – 1590) and Michael Praetorius (1571 – 1621) have established the basis of modern (western) musicological analysis in various formal frameworks. Moreover, as classical music covers more than 1,000 years of music repertoires, many concepts on these frameworks have varying semantics with time. For example, the concept of *Rondo* can invariably designate: (a) 14th century French songs structured as ABaAabAB; (b) Romantic movements of a piano sonata/concert structured as ABACADA; (c) 18th-19th century farewell airs sung by the Prima donna in operas. Considering that these challenges might be too broad and deep to be solved in a 3-year project, in Polifonia we address instead specific parts of this time-aware formalisation with certain constraints. First, we do this under the specific setting of **Semantic Web technologies and ontologies**, with well-defined expressivity and limits regarding what aspects of musical theories can be formalised. Second, we address these challenges with the basic tenet of **reusing existing bodies of knowledge** and making them useful for the specific goals of the project, instead of reinventing new ones, by e.g. proposing representations of the theories of Zarlino/Praetorius as OWL ontologies. Third, we plan on addressing the problem of **change of meaning through time by combining novel techniques in the Semantic Web**, such as temporal knowledge graph models and embeddings [11, 12], and **accompanying ontologies with high-quality documentation** in order to support the users of our tools at making a correct interpretation and use of time-sensitive concepts.

In this deliverable, we report on the advances made around the **Polifonia Ontology Network (PON)**, a set of networked ontologies designed with the specific purpose of identifying key resources for the separate description of a **musical object** and its **cultural and historical context.** To ensure a high degree of integration and reuse of the various components in the PON, we propose an ontology engineering methodology that focuses on the development and enrichment of various *Ontology Design Patterns* that adhere to the eXtreme Design (XD) methodology [13]. In this report, we describe the implementation of XD for the collected requirements, competency questions, and datasets of Polifonia; we describe the ontology modules of the PON that stem from the application of the XD methodology to such requirements; and we summarise a two-fold evaluation based on competency question validation by the Polifonia Music Expert Community (PMEC) and a use-case in developing a demonstration for the international event **Music & AI** organised by the widely known **Sónar**, an arts, design, and electronic and experimental music

[2] In this deliverable we refer to *music objects* as any representation that conveys the semantics and meaning of a musical work, such as performances, recordings, musical notations, scores, and so on.

festival founded in Barcelona that gathers yearly around 126k visitors. Additionally, we report on the first steps in using the resources developed in the PON to convert and annotate the Polifonia datasets into a first version of the Polifonia Knowledge Graph (PKG) in RDF (Resource Description Framework) [14] format.

The rest of the report is organised as follows. In Chapter 6 we survey related work in ontology engineering methodologies, ontologies for musical heritage, and available datasets. In Chapter 2 we describe eXtreme Design as our chosen engineering methodology, as well as its specific implementation for the particularities of Polifonia. In Chapter 3 we describe the PON, in particular its creation process using the methodology discussed in 2 and the resulting ontology modules. In Chapter 4 we propose an evaluation of the PON based on expert validation of competency questions and a use-case. In Chapter 5 we summarise our steps in using the PON to generate a preliminary version of the PKG, before we conclude in Chapter 7.

2 Methodology

2.1 eXtreme Design

The development of the Polifonia ontology network and Knowledge Graph (KG) relies on eXtreme Design. eXtreme Design (XD) has been selected as an ontology engineering methodology since it provides methodological support for incrementally addressing small sets of requirements and for minimizing the impact of changes in incremental releases, hence it's specifically useful for a project, as Polifonia, which aims at addressing a huge and diverse domain. Moreover, Polifonia's ontology engineers team has expertise in building ontologies and ontology networks based on XD. Finally, eXtreme Design has already been applied to the cultural heritage domain, which gave also input to some useful extensions [15].

eXtreme Design [16, 17] is an ontology engineering methodology that focuses on, and provides guidelines for, the reuse of ontology design patterns (ODPs). Ontology patterns provide modelling solutions to recurrent modelling problems: if reused as small template ontologies, they guarantee a high level of quality of the ontology, and favour its reusability [18]. Being XD an agile methodology, it is iterative, test-driven, and produces incremental releases, corresponding to different stages of the development. A collaborative development involves 4 teams: (i) a *customer team*, which provides the requirements that guide the design and testing steps; (ii) a *design team*, which models the ontologies and ODPs; (iii) a *testing team*, which takes care of testing and validation of the results; (iv) an *integration team*, which focuses on the integration of the produced components.

Requirements collection and competency questions. A preliminary step consists in collecting requirements from the *customers*, and in engineering them. In XD, requirements are collected in the form of *user stories*, which aim at describing by example the kind of facts that the ontology needs to represent, e.g. "*Tosca* was performed in Rome on 14 January 2000". One or more competency questions (CQs) [19], i.e. questions we want to answer against a knowledge graph, are derived from a generalisation of the user stories. For instance, from the previous story we could derive two CQs: "Where was a musical composition performed?" and "When was a musical composition performed?".

From CQs to pattern-based ontology engineering. At each iteration, a set of coherent CQs is selected, and possible existing solutions (ODPs) in ontologies or online catalogues of patterns[1], are analysed by matching the respective competency questions in order to find the most suitable one to be reused, or specialised, in the ontology under development. For instance, the CQs in the previous paragraph can be generalised in "Where and when a situation took place?", which

[1] Such as ontologydesignpatterns.org

can be matched to the ODP *TimeIndexedSituation*[2].

Testing and integration. In the context of the testing activities, CQs are translated into possible SPARQL queries, and the testing team verifies if the developed ontology can answer those queries. Instead, general constraints (e.g. “A time interval is disjoint with a place”) are tested by creating sample RDF triples, that are expected to provoke either consistency/coherence errors or inferences.

Integration tests focus on re-running all previously defined tests and running new integration tests on the whole ontology. The results of the tests are reported to the design team.

eXtreme Design has been applied to existing projects, including projects in the Cultural Heritage domain, such as those presented in [20, 21].

In the next two sections we show how we adapted eXtreme Design to Polifonia.

2.2 XD in Polifonia: stories, personas, and competency questions

In eXtreme Design, a story-based approach guides the collection of the project's requirements. As defined in the context of Polifonia, a *story* is a framework for *customers* to describe their needs, and is composed of four sections: the persona, the goal, the scenario and the competency questions (CQs). The *persona* is a research-based description of a typical user. It contains attributes such as name, age, occupation, and relevant characteristics of the person such as their knowledge and skills and their interests. The *goal* is a short textual description of the goal(s) that the persona needs to be addressed in the story. The goal(s) is(are) also represented by a short (maximum 5) list of keywords, to be provided by the customers. The *scenario* is a story describing how the persona's task/need/problem is solved before, during and after interaction with the resource/software/service being developed. Lastly, the *competency questions* (CQs) are question(s) the persona needs the resource/software/service to answer for satisfying their task/need/problem. Competency questions are the most crucial information that the ontology engineers retrieve from the stories. These CQs are direct requirements used for the modelling of modules that shape the ontology network. Nevertheless, personas and scenarios can still play an important role in ontology engineering in facilitating knowledge engineers a better understanding of the **provenance and original context** of competency questions (which are, as mentioned, the carriers of the requirements for ontology design). A full description of the roles of personas, scenarios, use-cases, requirements, etc. can be found in deliverable *D1.1 Roadmap and pilot requirements, 1st version* [1]

In cooperation with the expert's community, we have created the profile of 13 personas and formulated a total of 19 stories following guidelines described here[3]. In Figure 2.1 is shown an excerpt of one story involving the persona Carolina.

[2] http://ontologydesignpatterns.org/wiki/Submissions:TimeIndexedSituation
[3] https://github.com/polifonia-project/stories

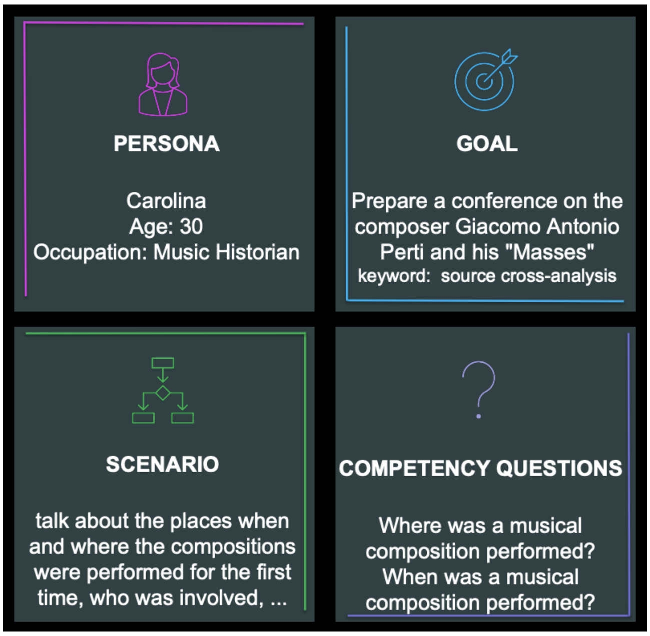

Figure 2.1: Excerpt of a story of the persona Carolina.

2.3 Analysis of requirements and setting of goals

All the competency questions, collected as described in the previous section, were individually analysed in order to identify any potential inconsistency that would create potential obstacles for the ontology design activities. In particular, after a preliminary analysis of the competency questions across all personas, several inconsistencies emerged.

1. It is common that a CQ is formulated in a way to include vague concepts and undefined terms. For instance, asserting whether two compositions are "*connected*" without providing the specific property on which the connection can be established (e.g. similar melodies, rhythm) does not provide enough context to the ontology design team. Without a more specific definition, it would be virtually impossible to formally express ontological requirements from these competency questions.
2. CQ can be complex or nested – entailing more than a single requirement as a result of nested logical operators articulating the question, or implicitly calling for a number of non-trivial operations/steps in order to be addressed. Before being further processed, such competency questions should be conceptually decomposed/simplified into a number of atomic CQs that can be associated to basic requirements. For instance, a prototypical CQ asserting "*How is track B connected to C? Can we conclude D?*" should be decomposed into: "*How is track B connected to C in terms of Y?*" (also note the contextual specification, as argued before), and "*Assuming Z, can we conclude D?*".
3. CQs may not be consistent with each other, as they originate from different personas, use cases (stories), and were not annotated/produced by the same XD practitioner. Indeed,

a cross-CQ analysis can only be operated at this stage, as any earlier intervention would potentially bias the interview process and the resulting collection of requirements.

In addition, we also found that CQs may include technical jargon and a level of musical detail that could not be fully understood by the ontology design team. In these circumstances, searching technical terms on the Web for clarification, or consulting domain-specific encyclopedic sources is not always a viable solution, as the context of such terms is usually defined/set within the CQ.

To address these issues, we designed a simple analytical framework where the critical points of each competency questions were explicitly highlighted in tabular format so as to bring them to the attention of story creators (the XD practitioners producing CQs from technical interviews) and domain experts. The reformulation of competency questions is indeed possible if those who originally contributed them can provide further information and approve changes proposed by the ontology design team to tackle their issues (1-3).

WELCOME story-creator or domain-expert (or both)

(*) Each persona, with his/her own stories, is analysed in a separate sheet in this document. Please, find yours!
(*) We encourage you to familiarise with our analytical framework (see the tables below) before jumping to your story.
(*) The sheet "Vocabulary and Tags" collects a list of domain-specific terms which will be annotated and described after the domain experts will clarify our points.

Each persona is analysed in relation to the competency questions (CQs) provided in the GitHub's stories, according to the following schema.

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the original story (on GitHub)							
							creator: John Patitucci
					domain experts: Jaco Pastorious, Victor Wooten (assigned by the expert team)		
Persona[story]-CQ[number]	Label of a proposed intervention.	Textual description of the original question, which can be complex, ambiguous or potentially not consistent for the ontology design team.	either design or visual	ID of a related CQ, potentially in another story.	A list of tags for ontology design.	Short pointers to the ambiguous parts of the CQ. You will find (at the bottom) a specific question that we will kindly ask you to address.	Whether the creator needs to approve a CQ replacement or a CQ split.
Pippo1-CQ1	-	How is instrument B connected to C? Can we conclude D?	design	Pluto1-CQ3	INSTRUMENT	Definition of "connection"	-
	A	How is instrument B connected to C in terms of Y?					x
	B	Can we conclude D?					approved
Pippo1-CQ2	-	How composer C is realated to composer D when bla bla?	design		COMPOSER	-	
	R	Is composer C the Z of composer D?		-			x

At the bottom of each sheet, and depending on the ambiguities identified before, you will find a table with our questions. Please check if a question requires your attention (Expert and Creator flags).
NB: when replying to our questions, we would appreaciate if you could use basic, simple and non-ambiguous terms, as if you are explaining a certain concept to a kid. The simpler, the better!
Also, please avoid changing the style of the document: if your explanation-answer does not fit the size of a cell, just click on the "wrap" button (next to the text alignment button). Thanks!

Figure 2.2: Overview of the analytical framework for CQ resolution: a simple iterative methodology for improving competency questions from expert feedback.

As shown in Figures 2.2, the technical analysis of CQs is focused on the stories of each persona, and outlines their context (a list of keywords/tags characterising the main concepts/entities involved), the type of requirements they entail (whether visual or design-oriented), their critical issues (potential ambiguities and inconsistencies that need to be resolved), and possible relationships with other competency questions (also belonging to other personas). The table also includes a column called "*breakdown*" together with additional CQs (those denoted by capital letters in the former column) that correspond to the resolution of an original CQ suggested by the ontology design team after the analysis. When prompted to clarify ambiguities and inconsistencies – using the interaction panel in Figure 2.3, a story creator is also asked to verify the suggested breakdown of complex/nested competency questions. This last operation can be either approved, amended, or rejected with proper justifications[4].

[4] In the latter case, the ontology design team will repeat the analysis – using the additional material provided by the

Q/A Logbook					
CQ number(s)	**Type**	**Argument**	**Expert**	**Creator**	
3, 4 (the CQs' number)	Q	What do you mean by X?		**needed**	*Example of a question that only needs to be addressed by the story-creator*
	A	What I meant by X is Y, for example Z.		John	*Example of a story-creator's answer.*
10	Q	How can we formally define B?	**needed**		*Example of a question that only needs to be addressed by a domain expert.*
	A	B is defined as L, M when etc.	Jaco		*Example of an expert's reply.*
11	Q	The concept of connection needs to be clarified.	**nedeed**	**needed**	*Example of a question that requires the attention of both.*
	A (E)	The term connection is usually used to identify bla bla bla	Victor		*The expert replies.*
	A (C)	Connection was intended as bla bla bla		John	*The creator replies.*

Figure 2.3: Summary of the action points expected from either, or both, story creators and domain experts – in relation to a single or a group of competency questions.

In case the ambiguities are relatively technical, and cannot be fully explained and resolved by the story creator, a domain expert (e.g. a musicologist) is assigned to the corresponding CQ. In doing so, they can provide relevant material and simple explanations (c.f. Figure 2.3) useful to clarify the terminology and solve the ambiguity. The same interactive flow with domain experts is also considered when the ontology design team encounters highly technical terminology and needs further clarification of such terms within the context of their CQs.

2.3.1 Reusing existing resources

The ontological needs of Polifonia are diverse, and require describing cultural heritage objects from different contexts, formats, historical periods, and levels of detail. Existing ontologies, outlined in the previous subsection, do not address all the Polifonia requirements collected so far, thus new modelling solutions are needed. However, there are relevant entities and fragments from these and other (e.g. DOLCE [22], ontology design patterns [17]) ontologies with valuable, overlapping knowledge, and thus with a great potential for reuse here.

In Polifonia, the decision on reusing existing ontological resources , instead of creating them anew, depends on the project requirements, as provided by the pilots in [1]. In other words, the essential criteria for reusing existing ontological resources (properties, classes, ontology fragments) is whether such a reuse will contribute to the ontology answering one or more competency questions.

After deciding whether reuse is appropriate, we establish criteria for **choosing between a direct and an indirect reuse** [23, 24]. *Direct reuse* consists in either importing the whole external ontology, or including selected ontology terms, with their original URIs, in a new ontology. Contrarily, *indirect reuse* means that relevant entities and patterns from external ontologies are used only as templates: they are reproduced (and possibly extended) in the local ontology, thus decreasing the dependency on external ontologies. Both approaches have advantages and disadvantages: direct reuse can be faster and less laborious (as the engineering work is directly imported), and can increase the interoperability and semantic capabilities of ontologies. On the other hand, the evolution of the reused ontologies is outside the control of the ontology engineers that are reusing them: changes in the original ontologies, or ontologies that are no longer maintained, may cause inconsistencies and instability in the local ontology. Therefore, as in Polifonia it is essential to

story creator, and iteratively repeat the procedure for the verification of any suggested CQ modification.

guarantee control and independence of the ontologies to answer the pilot's competency questions [1], we adopt an indirect reuse. To facilitate further integration and to guarantee a higher degree of interoperability than provided by simple indirect reuse, we aim at providing alignment axioms (e.g. through the use of `rdfs:subClassOf` and `rdfs:subPropertyOf`) at a later stage.

Therefore, the process of materialising alignments between the entities defined in the Polifonia ON and those reused from existing ontologies and patterns is ongoing. Examples of reused ODPs are *TimeIndexedSituation*[5] – which has been specialised for representing musical performances – and *Time Interval*[6], while the ontology model for the recording process and its sessions is reusing (and will be aligned to) concepts from the Music Ontology.

[5] http://ontologydesignpatterns.org/wiki/Submissions:TimeIndexedSituation
[6] http://ontologydesignpatterns.org/wiki/Submissions:TimeInterval

3 Ontology Network

In this Chapter we summarise the work produced by applying the XD methodology to the personas, stories, and requirements derived from the WP1 Pilots [1] in order to create a first version of the **Polifonia Ontology Network**. We describe in detail the curation process that followed from collecting, analysing and normalising requirements into competency questions that can be modelled and satisfactorily answered through various ontology modules. This collection, analysis and normalisation of competency questions was done in collaboration with a community of domain experts within Polifonia. This led to the identification and modelling of 10 ontology modules: the *Core module*, the *Musical Performance module*, the *Musical Composition* module, the *Musical Feature module*, the *Source module*, the *Instrument module*, the *Comparative Measure module*, the *Music Emotion module*, the *Bell module*, and the *Metadata* module.

3.1 From CQs to Ontology Engineering

Using our analytical framework described in Section 2.3, most of the issues emerged from the competency questions were solved thanks to the active involvement of story creators and domain experts. A detailed overview of the analytical efforts and outcomes – for all current personas and a selection of CQs, is reported in Appendix 7 for further analysis.

The improved competency questions resulting from this collaborative and iterative interaction, could then be translated in clear, atomic and consistent ontological requirements. Considering the wide diversity of CQs – ranging from general events to musicological interpretations of specific passages in compositions, the first step was to cluster/categorise them in thematic groups. This process was driven by the description of each CQ's context – a list of meaningful tags and keywords describing the main entities and concepts explicitly mentioned or implicitly related in/to a CQ. All these descriptors were collected during the analysis stage, as outlined in Section 2.3. Overall, more than 110 unique keywords were identified, which in turn were categorised in 10 thematic groups. The results of this taxonomy directly informed the design of the architecture of the Polifonia Ontology Network.

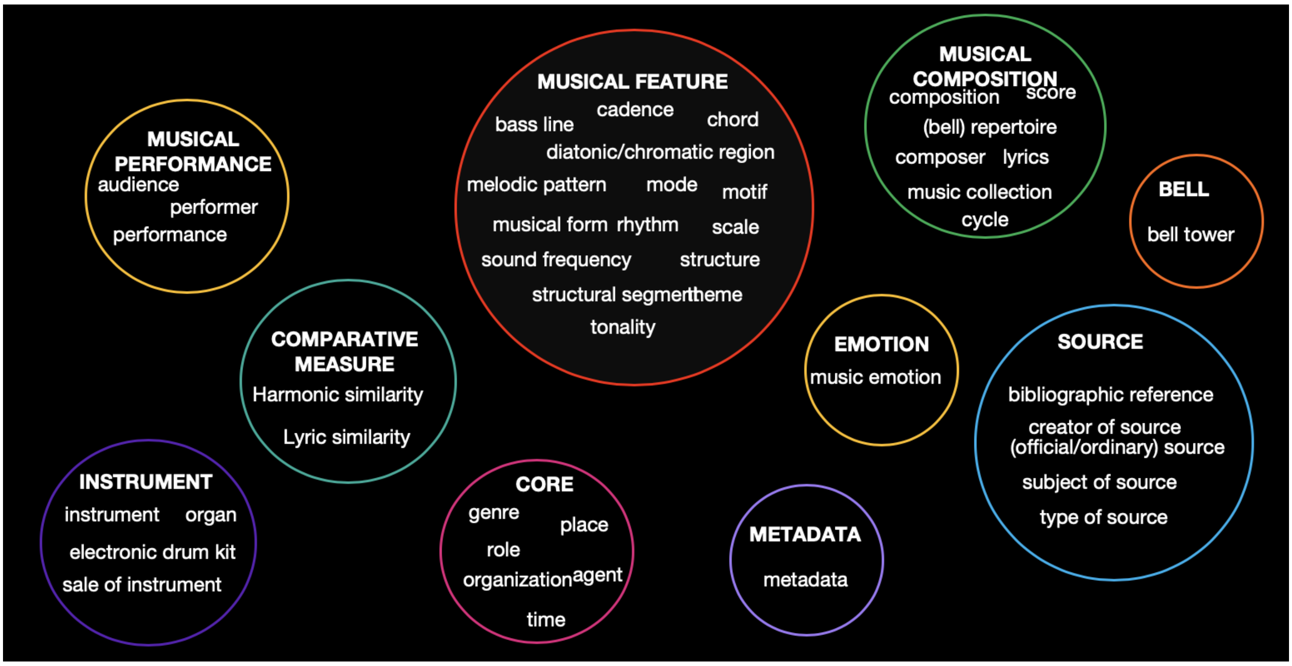

Figure 3.1: Taxonomy of tags collected from the analysis of CQs into thematic groups.

3.2 Description of the Ontology Network

The Polifonia Ontology Network, with all links available in the dedicated repository[1], is currently composed of 11 ontology modules. A root ontology module (FULL) imports all thematic ontology modules of the ontology network. Thematic and top-level ontology modules are stored elsewhere in their own repositories, and linked here. In figure 3.2 we present the Polifonia Ontology Network, composed by the modules that are described in the following subsections. Table 3.1 contains the links to the repositories storing the modules, their prefixes – that are used in the following figures, and references to the stories from which each module takes its requirements. Actual stories can be found in the *stories*[2] GitHub repository.

3.2.1 The Core Module

This module models general-purpose concepts and relationships (e.g. place, time, classification, situation) orthogonal to the whole Polifonia ontology network, and is imported by all the other thematic ontology modules of the network.

Figure 3.3 shows the main classes and relations for modelling places and place-related information. A `core:PhysicalSite` is an area, a structure (e.g. a building) or a group of structures where something was, is or will be located. It is a physical site that situ-

[1] https://github.com/polifonia-project/ontology-network
[2] https://github.com/polifonia-project/stories

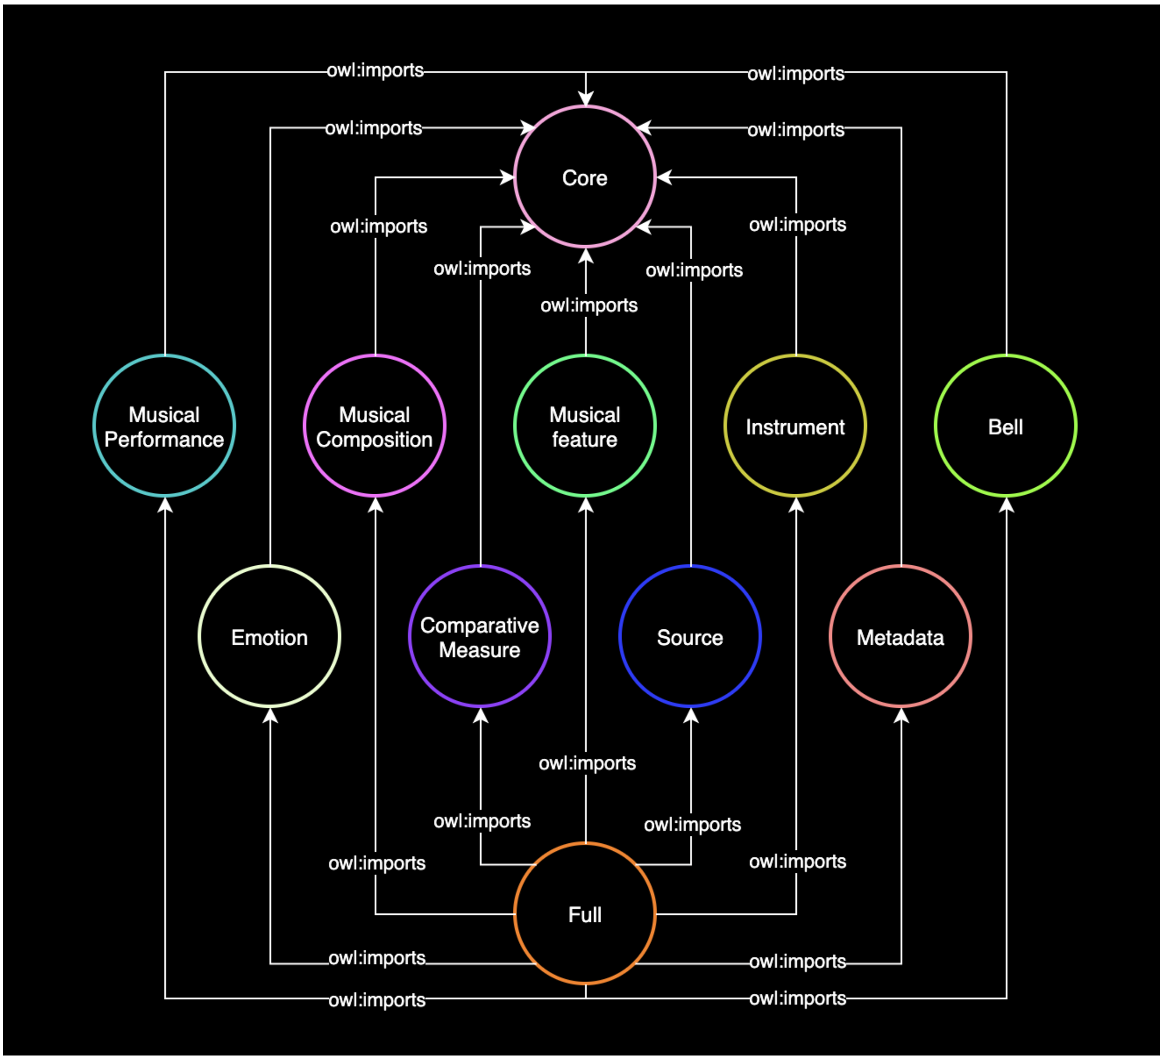

Figure 3.2: Polifonia Ontology Network.

ates something. The property `core:hasAddress` relates a physical site to its address, which is in turn related to its `core:AddressComponents`, such as the Country and the City. Any `core:GeographicalFeature` has a `core:Geometry`, defining latitude, longitude and altitude. A `core:Person` is associated with her birth place with the property `core:hasBirthPlace`. As in Figure 3.4, we represent a `core:TimeInterval` by associating it with its start and end time.

3.2.2 The Musical Performance Module

This module models concepts and relationships for representing musical performances and music-related events. It already addresses a set of competency questions related to `mp:MusicalPerformance`s, modelled as time-indexed situations, including the place where,

Ontology Module	Prefix	Stories	Repository
Full	full:	n/a.	~/ontology-network
Core	core:	n/a.	~/core
Musical Composition	mc:	Carolina#1 - Sources cross analysis	~/musical-composition
Musical Performance	mp:	Carolina#1 - Sources cross analysis Sonia#1 - Exploration mode Sonia#2 - ShuffleMode	~/musical-performance
Musical Feature	mf:	Anna#1 - HearingMusic Keoma#1 - Restoration And Sound Practices Mark#1 - Folk Music Sethus#1 - Conflicting Theoretical Interpretations Sonia#1 - ExplorationMode Sophia#1 - Musicians and their environment	~/musical-feature
Music Emotion	me:	Ortenz#1 - Music And Childhood Sophia#1 - Musicians and their environment	~/music-emotion
Bell	bell:	Keoma#1 - Restoration and sound practices	~/bell
Source	src:	Carolina#1 - Sources cross analysis Keoma#1 - Restoration and sound practices Ortenz#1 - Music and childhood Sonia#1 - Exploration mode Sophia#1 - Musicians and their environment Sophia#2 - Origins and form	~/source
Instrument	inst:	Carolina#1 - Sources cross analysis David#1 - MusicHistorian Mark#2 - Dutch Organs Sophia#1 - Musicians and their environment	~/instrument
Comparative Measure	cm:	Sonia#1 - Exploration mode Sonia#2 - ShuffleMode	~/comparative-measure/
Metadata	md:	n/a.	~/metadata

Table 3.1: Overview of the modules in the Polifonia Ontology Network, together with the specific personas/stories addressed by each of them. In the last column, the first name denotes the persona, and the number next to the # symbol identifies the story, followed by its title. All the URL of the corresponding repositories is relative, with respect to the Polifonia codebase at https://github.com/polifonia-project/.

and time interval when, a musical composition has been performed (for the first time), and the performers (person, choir, group, etc.) involved with different roles, such as singer or instrument player (see Figure 3.5). A musical performance can be composed of different musical performances, and this can be represented with the property `core:hasPart` between musical performances.

Moreover, the current version includes classes and relationships for representing the activity of recording a musical performance (see Figure 3.6), with possible sessions (e.g. mixing, recording)

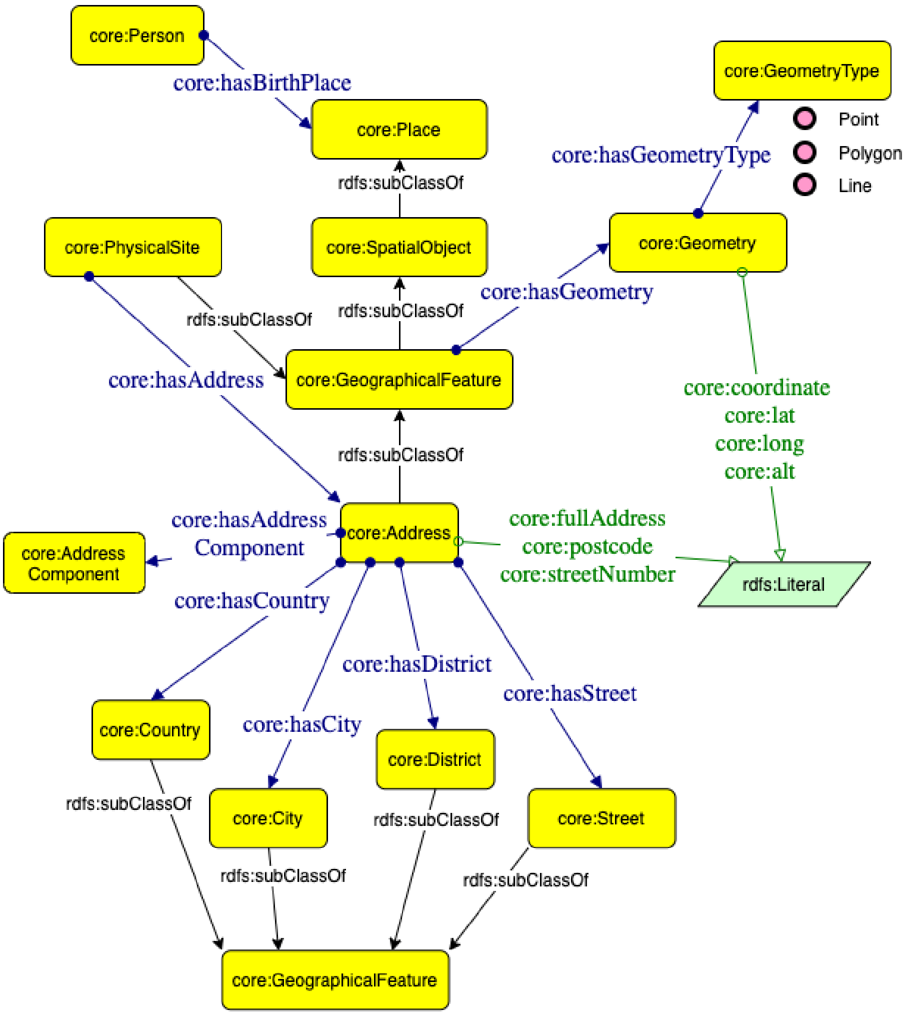

Figure 3.3: The model for places.

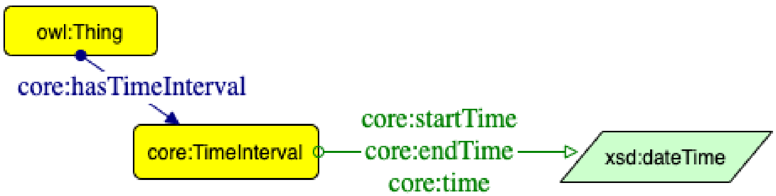

Figure 3.4: The model for time intervals.

and the related recording produced as a result of this activity, with its attributes, such as the title.

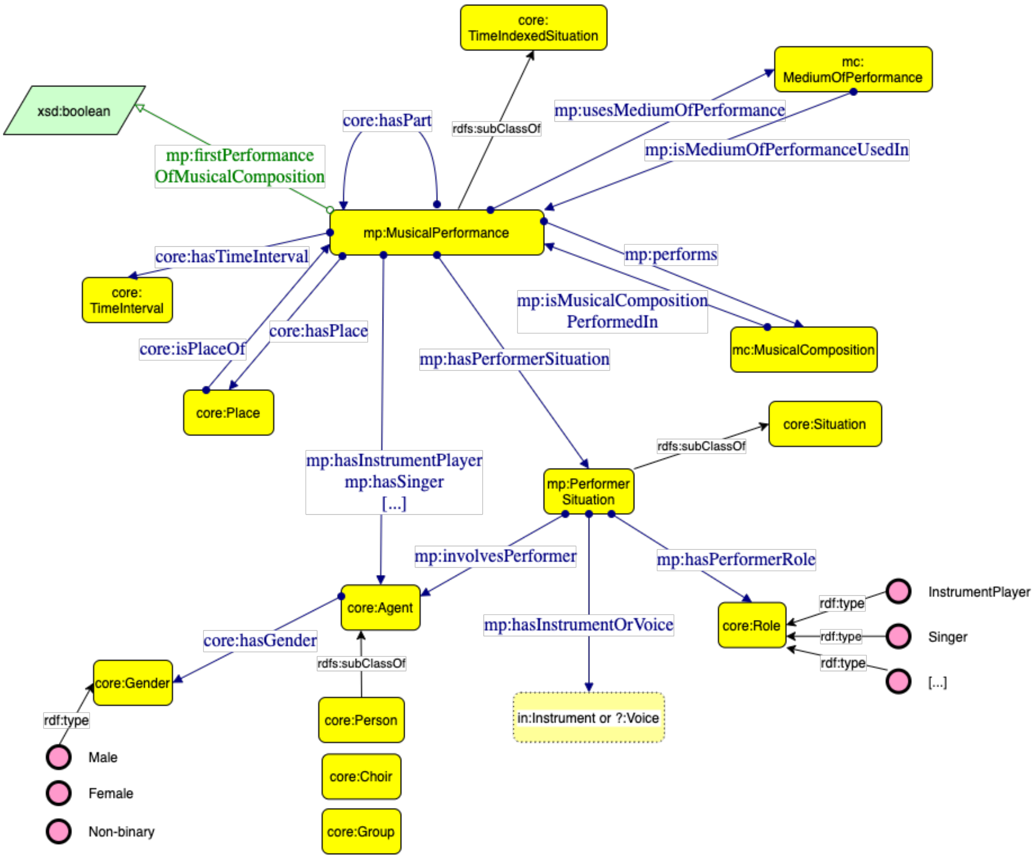

Figure 3.5: Model for musical performances.

3.2.3 The Musical Composition Module

This module models musical compositions and related entities, such as its composer(s), with the property `mc:hasComposer`, and its medium of performance (see in Figure 3.7 the class `mc:MediumOfPerformance`), intended as the voices and instruments needed to perform a piece of music. There are compositions with: a) only instruments; b) only voices; c) voices and instruments. The medium of performance adopted during a specific musical performance can be different from the one that has been defined by the composer. Moreover, this module represents the lyrics that can be associated with a musical composition.

3.2.4 The Musical Feature Module

The musical feature module should provide a comprehensive and sound schema to describe a musical object (a score, an audio track, etc.) in regard to the musical properties that can be objectively attributed or subjectively identified from it. Hence, musical properties can be assigned to musical material — the whole composition-performance, or a part of it — as a result of an interpretation. In addition, it is also worth distinguishing between a musical feature that has

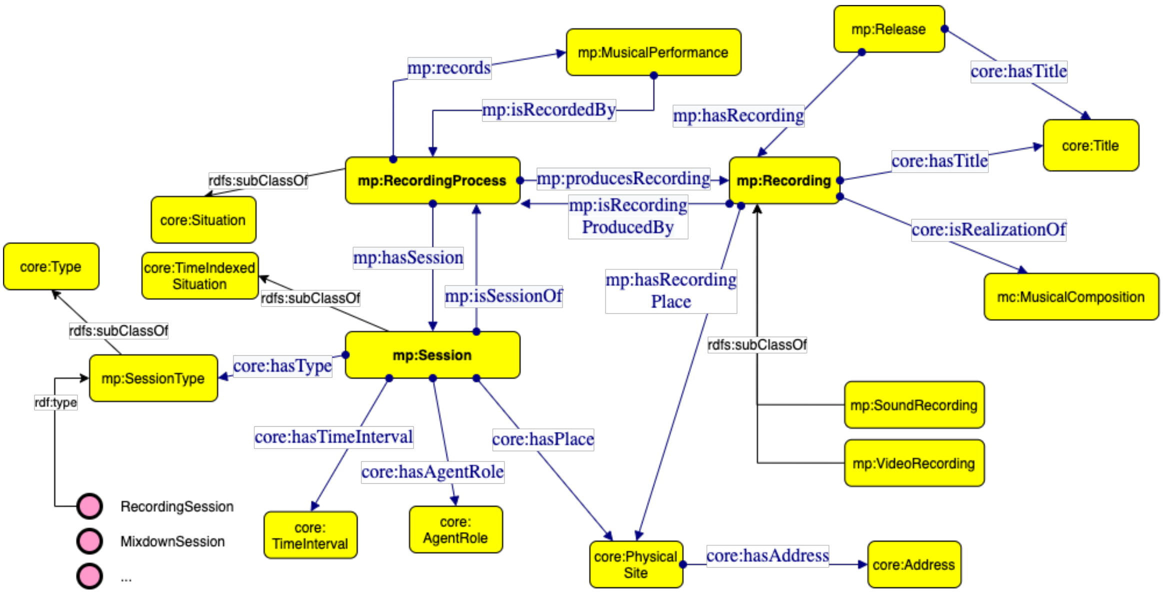

Figure 3.6: Model for recordings.

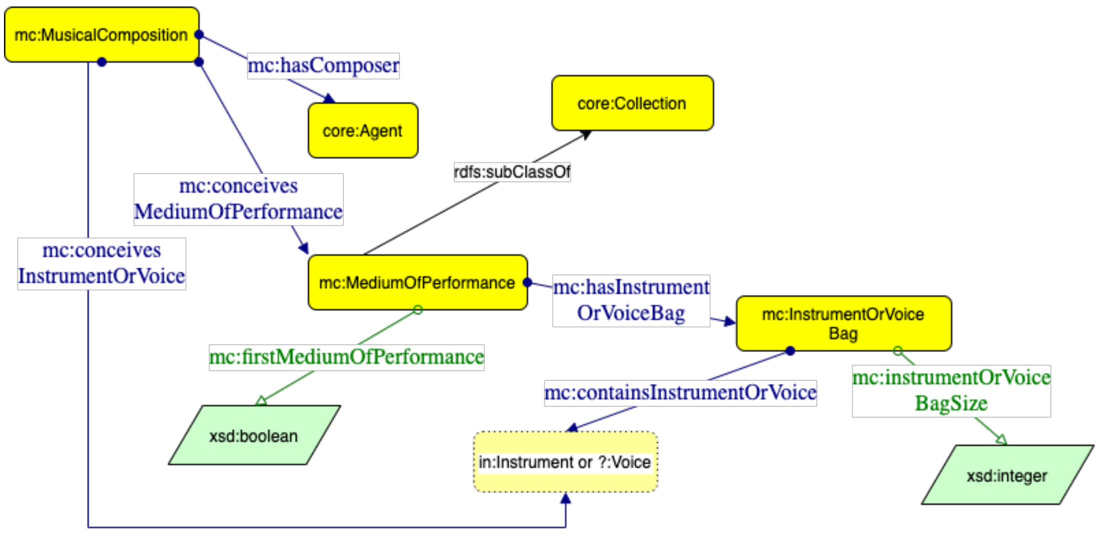

Figure 3.7: Model for musical compositions.

been assigned manually (human annotator) or inferred by means of an algorithmic procedure (computational annotator) — a distinction that we attempt to relate to provenance. As several concepts need to be modelled in order to enable an exhaustive description of musical material, this module can grow considerably and further subdivisions (sub-modules) could be identified from it. For the moment, we will condense and accumulate everything in this module.

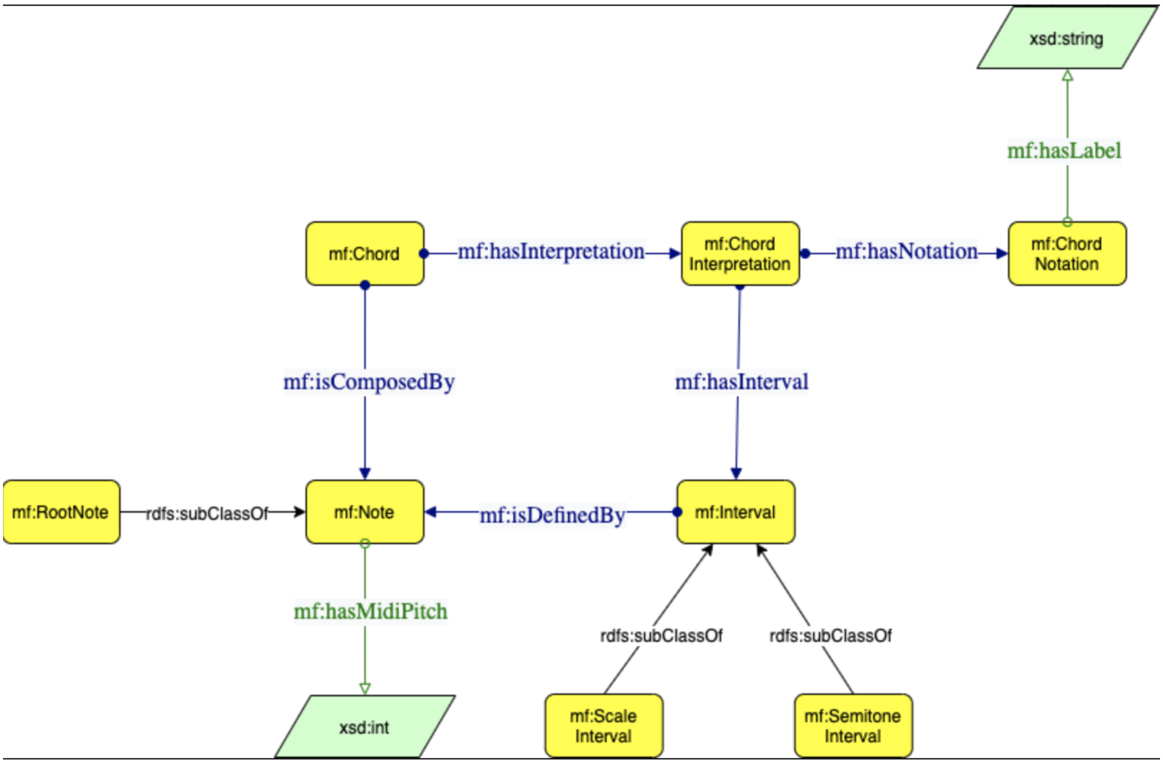

Figure 3.8: Model for chords as musical features.

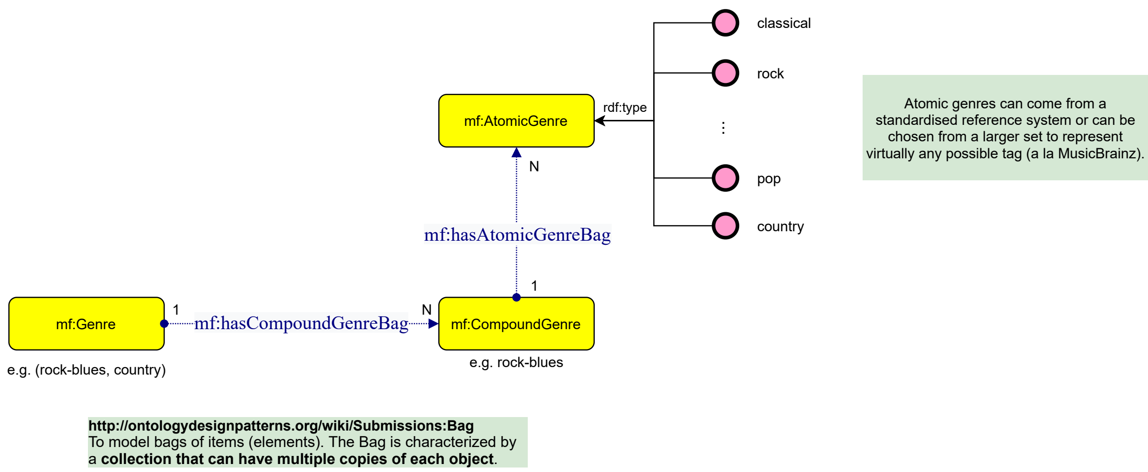

Figure 3.9: Model for genre as a musical feature.

3.2.4.1 TONALITIES ontologies

Several models relating to historical and present-day modal and tonal theory are being developed as part of the TONALITIES pilot. These models are the starting point for classifying musical works into modal-tonal categories on the basis of an in-depth understanding of their inner or-

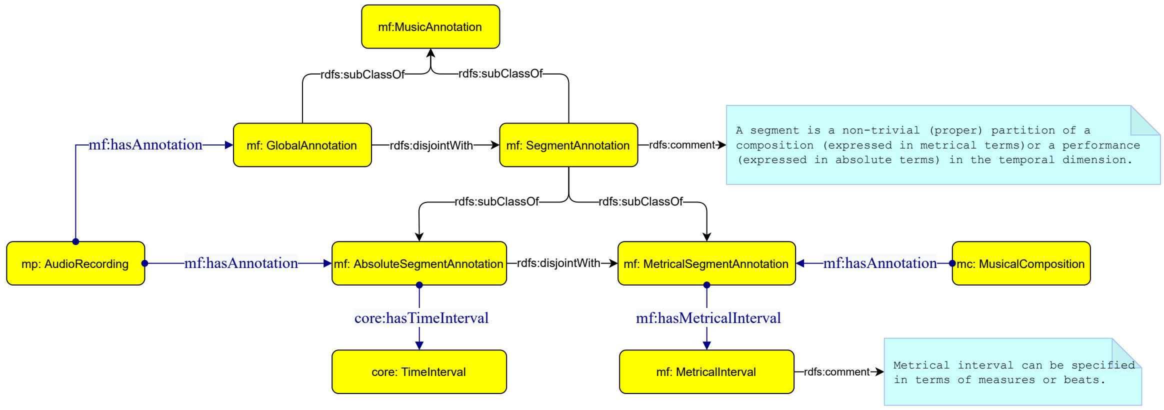

Figure 3.10: Model for other musical features.

ganisation. The modelling of different theories relating to the same knowledge domain - i.e. modal-tonal organisation - aims to avoid the limitation to a single "standard" and "global" theoretical perspective, arbitrarily chosen and applied mechanically despite its limited validity (belonging to a historical period, restricted to certain repertoires, conveying certain ideologies, etc.). The strategy adopted, outlined in the Sethus story, consists in letting the user apply different models, corresponding to distinct theoretical viewpoints, to observe how these perspectives affect the modal-tonal understanding of the same musical work. Beyond, this multi-faceted approach allows for the confrontation of models in order to grasp how theoretical thoughts evolve over time and space. The alignment of these ontologies then shows how concepts – for example mode, ionian on D, transposition, cadence, chord, etc. – evolve while the words to denote them remain the same and, conversely, how certain concepts can remain identical while their names differ.

At the present stage, two historical theories are available: Zarlino 1558 and Praetorius 1619 (with about 2500 axioms, 500 classes and 100 object properties each). The Praetorius model has been presented [25] and applied [26, 27] in former works. Initial tests have been undertaken during the last weeks to infer the mode of individual works from these models using a reasoner (HermiT). Furthermore, a set of functions has been developed to allow the comparison of the models, both at the class level and at the property level and to measure their differences and overlaps (see Figure 3.11). In the coming months, it is planned to test the two models on larger corpora, to extend the modelling to contemporary modal-tonal theories, and to make these models available through a dedicated interface.

3.2.5 The Source Module

The Source module models concepts and relationships for representing sources of (music-related) information. It aims to provide information about the *creator* of the source, the *type* of the source, the *time* and *place* when and where it was created, the *context of production* and *context of usage*, the *subject*, the *goal* and the *source credibility*. Currently, the conceptualisation of the module is based on bibliographical sources only, but we are considering extending to

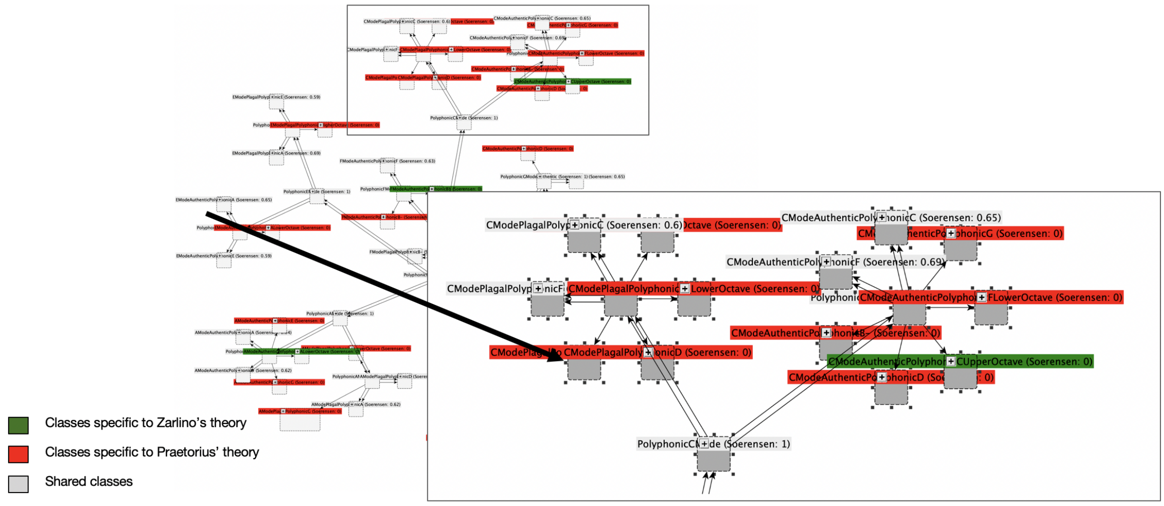

Figure 3.11: Models for the historical theories of Zarlino and Praetorius.

multimedia sources (image, sound, video) as well. In this case, we would also acknowledge new types of information that need to be added to the module, for instance *file format description*, *encoding*, *licence*, *copyright*, *language*, etc.

While modelling we have also identified a possible pattern describing source production, that can be applied for different types of sources (see Figure 3.12). This pattern covers information about the source itself, the *time interval* when it was produced, the *place* where it was produced and the *agents* that participated in the production of the source with their respective roles.

3.2.6 The Instrument Module

The Instrument module is created to describe musical instruments, their categorisation into instrument families and groups, and their technical properties. Considering that an important part of the this module is the representation of the categorisation of the instruments, with the help of domain experts we have surveyed several reputable categorisations that are described below.

The categorisations that we have surveyed are Hornbostel–Sachs[3], Musical Instrument Museums Online[4] (MIMO), Virtual Encyclopedia of World Instruments[5] (VEWI), Oxford Music Online, Music Instrument classification by Wikipedia[6] and also categorisations based on datasets like MusicBrainz[7]. Hornbostel–Sachs is a system of musical instrument classification devised by Erich Moritz von Hornbostel and Curt Sachs. It is the most widely used system for classifying musical instruments by ethnomusicologis and organologists. The categorization in MIMO is an updated

[3] https://en.wikipedia.org/wiki/Hornbostel%E2%80%93Sachs
[4] https://mimo-international.com/MIMO/instrument-families.aspx
[5] http://www.instrumentsdumonde.fr/famille/familles.html
[6] https://en.wikipedia.org/wiki/Musical_instrument_classification
[7] https://musicbrainz.org/instruments

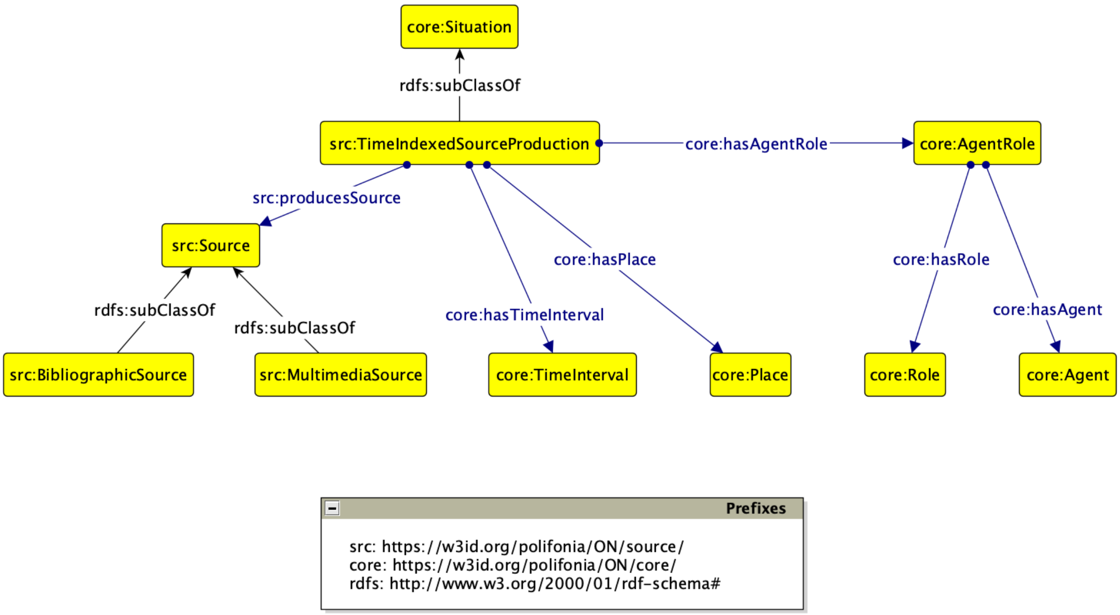

Figure 3.12: Model for source production.

version of the Hornbostel-Sachs categorization system. The categories of the instruments are:

- Wind instruments
- Percussion instruments
- Stringed instruments
- Keyboard instruments
- Mechanical instruments
- Electronic instruments
- Mirlitons
- Other instruments

While, in the VEWI, the categorisation is more general and it includes the following categories: woodwind, keyboard, string, brass and percussion instruments. Instruments are also distinguished by the way the sound is produced. Thus, we find cordophones (strings), membranophones (certain percussions), idiophones (the sound is produced by their material), aerophones (brass and woodwinds) and electrophones (digital keyboards for example). In addition, the Oxford Music Online introduces a new category of instruments that is corporeal techniques (body sounds) that include chest beating, finger snapping, overtone singing, tongue clicking, ululation, whispering, singing, etc. This new category provides argumentation that it is acceptable to include the human voice in the instrument module. Moreover, we have investigated related existing ontologies [28] that describe instruments. The problem with these ontologies is the fact that there are no other properties describing instruments besides the categorisation, which in consequence classifies the ontology more like a taxonomy.

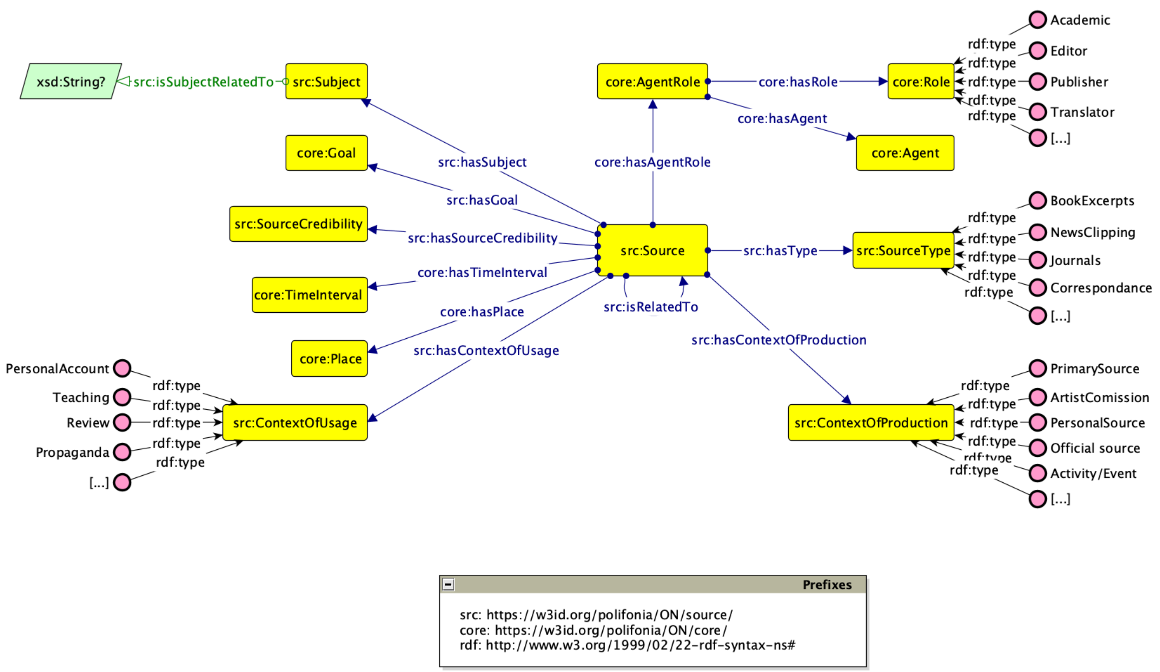

Figure 3.13: Model for sources.

Regarding the module itself, the requirements tackle the world of instruments in different perspectives and in order to make this module more inclusive we have created a new story, named Re-orchestration[8], for Sophia which deals with the categorisation of the instruments. For the modelling of the module, we have reused and specialised the `LinnaeanTaxonomy`[9] pattern which is content ontology design pattern represents the layered classification of the animal kingdom. Even though the domain where this pattern is originally applied to is different from the instrument domain, it can easily be adapted since the core idea of the pattern is classification of objects based on hierarchy. Therefore, we have by specialising this pattern, we have identified the classes of the categorisation that are `instrument family`, `instrument group` and `instrument`. Furthermore, in the module we have represented technical information of instruments such as `timbre`, `pitch`, `dynamic range`, `notation`, and historic information such as `builder`, `material inventor`, `place` and `time interval` of invention/production.

3.2.7 The Comparative Measure Module

The Comparative Measure models observations and measures applied to comparative analysis, such as similarities. The module shows great potential for representing novel computed information such as harmonic similarity, complete lyric similarity and line lyric similarity between

[8] https://github.com/polifonia-project/stories/blob/main/Sophia:%20Musicologist/Sophia%233_Reorchestration.md
[9] http://ontologydesignpatterns.org/wiki/Submissions:LinnaeanTaxonomy

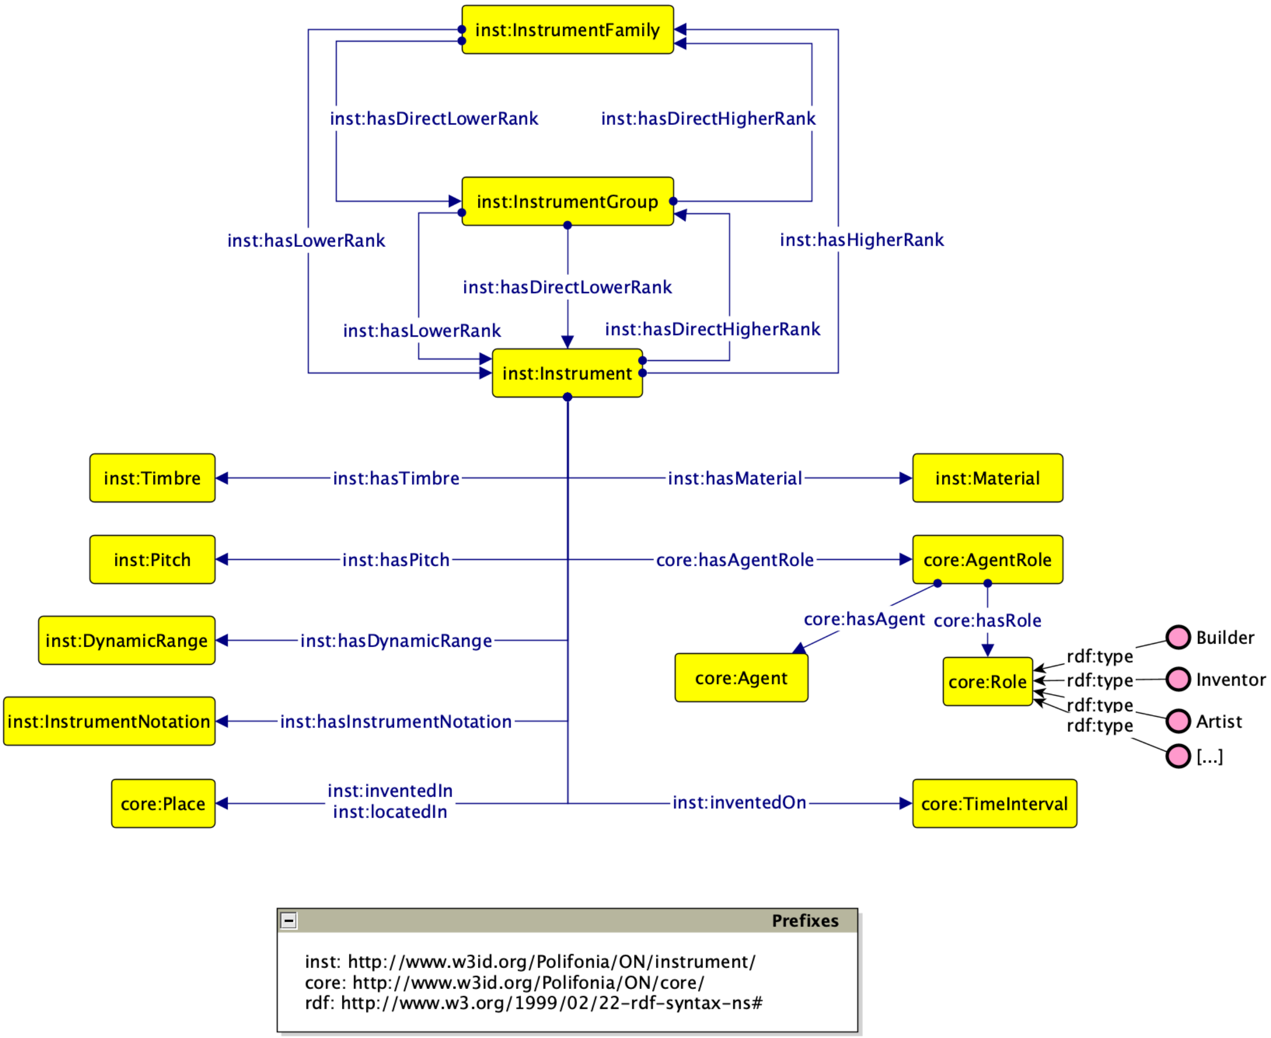

Figure 3.14: Model for instruments.

two musical pieces, source credibility, etc. Thus far, we have defined a general pattern for this module which can be applied for distinct types of similarities. In figure 3.15 we introduce the specialisation of the pattern to represent the harmonic similarities between recordings of two musical pieces. Common information that is presented in all types of similarities are the `similarity function`, `encoding`, `similarity score`, `recording`s that are involved in the similarity computation. While in the case of the harmonic similarity, there is the `time interval` that is specific to this specialisation because it represents the exact time interval where a harmonic similarity is detected between two musical pieces.

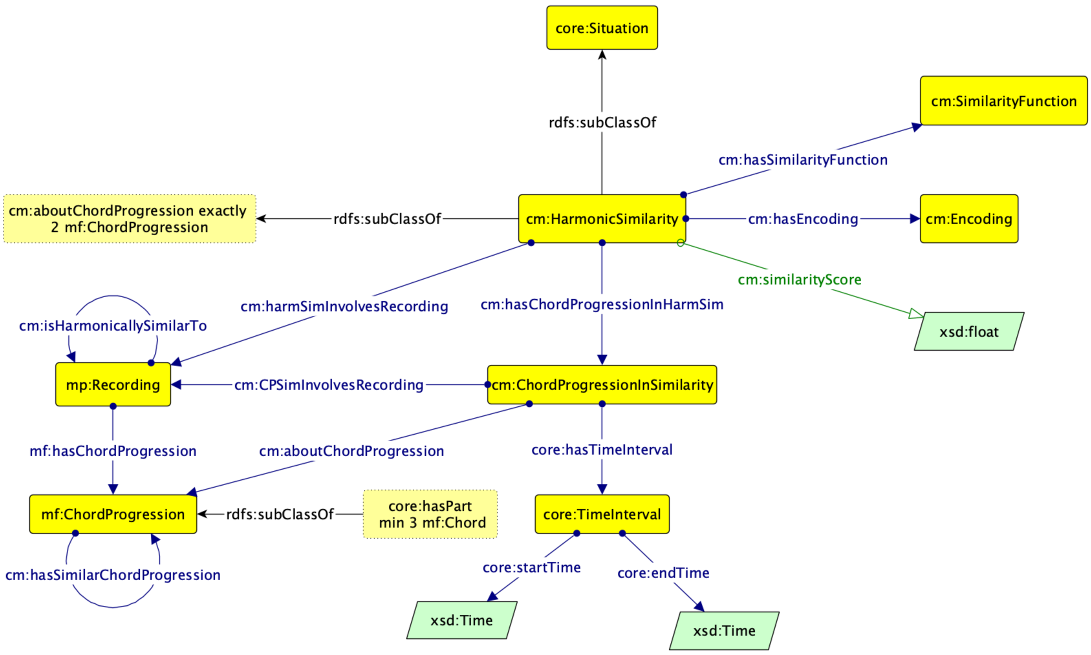

Figure 3.15: Model for harmonic similarity.

3.2.8 The Music Emotion Module

The music emotion module should provide (a first version can be found in Figure 3.16) a comprehensive and sound schema to describe emotions both perceived and induced by a composition in regard to the musical features. Hence, emotions can be assigned to musical material — the whole composition-performance, or a part of it — as a result of an interpretation. Before establishing whether a connection between emotion and interpretation is needed, there is a need to further understand which kinds of data (EEG signals, melodic patterns, facial expressions, verbal or textual sources) each of these types of emotion have as “evidence”.

3.2.9 The Bell Module

The Bell module aims at modelling the historical bells heritage, which includes both tangible heritage (bells, bell towers) and intangible heritage (performing practices, oral transmission of knowledge). The main entities that will be represented in this module are:

- the bell, with its physical characteristics (e.g. material(s) and dimensions) and historical information;
- the bell tower, which is a tower that contains one or more bells to enable them to be heard at a distance when rung, with physical characteristics and location information;
- the performing practices, that include gestures, instruments and body parts involved for ringing one (or more than one) bell;

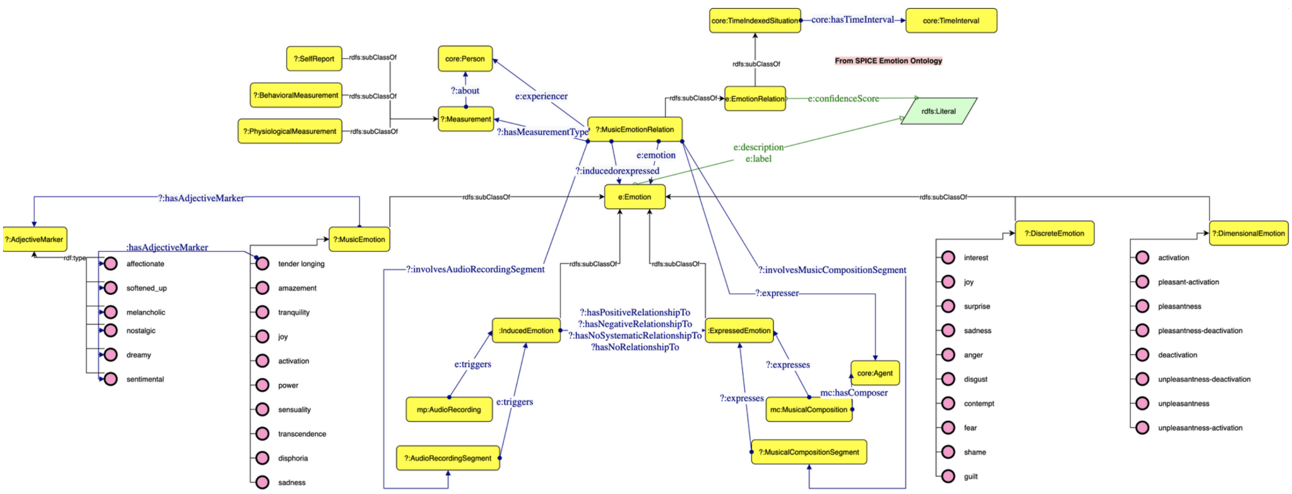

Figure 3.16: Model for music emotions.

- the bell-ringer and bell-ringing repertoires, as collections of pieces that a specific bell-ringing community traditionally performs.

3.2.10 The Metadata Module

To conclude, it is worth noticing that not only the Polifonia ON permits to model musical objects and their connections to cultural heritage artifacts, but it is expressive enough to provide a common standard for publishing and sharing music datasets on the Web. As argued in Section **??**, the lack of standards for releasing music data is indeed an open problem in the field of music information retrieval (MIR), with direct consequences on the final use of music datasets despite the considerable value of these resources. By leveraging the ontological ecosystem we are currently developing, Polifonia ON's classes and relationships will be reused to formally describe the collection and content metadata of arbitrary music datasets, together with their annotations (c.f. Section **??**). Although this particular use case does not currently address any competency questions explicitly, a metadata module would play a crucial role for the integration of MIR datasets – one of the main goals of the INTERLINK pilot. In fact, the ability to integrate music datasets would consistently contribute to the inclusion of data in the resulting KG, thereby providing diverse musical material and especially – expert annotations of musical features (chord progressions, local and global keys, structural segments, musical patterns, etc.) that would otherwise be hard to find or retrieve elsewhere with such a level of specificity and quality.

From a technical perspective, the metadata module will make extensive use of the other modules in the ON – by importing relevant classes and relations from them, and complement the latter with additional constructs needed to express further concepts at the metadata level (e.g. license attribution, curators, etc.) as well as relating them in the context of music collections (e.g a musical feature as an annotation type rather than an actual annotation). Overall, the musical performance and composition modules will be central to describe the content metadata of a dataset, whereas the musical feature module will serve classes and relationship to formally express the specific annotations a collection may provide. In this way, music datasets can already be connected at both the metadata and the annotation level – ready to be integrated within the same data infrastructure.

3.3 Reproducibility and FAIRness

As outlined in the beginning of the chapter, the development of each module in the Polifonia Ontology Network has already started to comply with open source best practises and standards. More precisely, although the actual release of the Polifonia Ontology Network is out of the scope of this report (and will be covered in future deliverables), each ontology module is currently maintained in a publicly available repository on GitHub (c.f. Table 3.1) and managed by one or more ontology designers. The modular design and the open source nature of the ontology network promotes both the reuse of components, as well as the creation of a community actively involved in the revision, improvement, and extension of the ontology modules. Not only this allows to continuously collect expert feedback within Polifonia and from the research community to a larger extent, but it will also simplify the implementation of the Data Management Plan when data from the pilots will be used to instantiate the Polifonia Knowledge Graph (c.f. Section 5). The Polifonia Knowledge Graph will adhere to the data publishing paradigm of Linked Data [3], and therefore be compliant with well-known data publishing practices like FAIR [2] (meaning data will be provided in a way that is Findable, Accessible, Interoperable and Reusable). This includes the fact that provenance information will always be included at all stages to trace the origin of data from each pilot and from each source. Since the PKG will be mostly implemented via Semantic Web technologies and Linked Data [3], provenance information will adhere to well-known standards such as W3C PROV [29] and nanopublications [30].

In sum, as the automatic methods developed for ontology testing and ontology validation will also be publicly released, the entire workflow – from the design and the evaluation of the ontologies, to the creation of the Polifonia Knowledge Graphs, will be entirely reproducible.

4 Evaluation

We follow well-established methodologies in the Semantic Web for ontology evaluation [31], which concerns the assessment of ontologies with respect to their usability, domain coverage, etc; and ontology testing [32]. To evaluate the modules of the Polifonia Ontology Network shown in the previous Chapter we propose a two-fold evaluation. First, we rely on the testing methodology of XD [13] and run tests on the generated competency questions (see Chapter 3) to assess whether the ON can provide valid answers to them or not. We do this with two important considerations: (i) that we aim at automating this testing and validation process of CQs; and (ii) that these CQ validation must be done with the expert knowledge of the Polifonia music expert community (PMEC). Second, we describe a use-case based evaluation, in which we used a part of the ON as a modelling tool to exchange and integrate datasets for the demo the project showed at the AI & Music festival at the Sónar International Music Festival on 27-28 October 2021.

It is important to consider that these are only the first evaluation efforts on the ON, and that the project has further tasks and deliverables that aim specifically at providing robust evidence for this evaluation; concretely (and not limited to):

- Deliverable *D2.4: Methods for interlinking knowledge graphs, 1st version* due on M18
- Deliverable *D2.7: Ontology testing and evaluation report, 1st version*, due on M18
- Deliverable *D2.2: Ontology-based knowledge graphs for music objects, 2nd version* (next version of this deliverable) due on M30

4.1 CQ driven evaluation

The eXtreme Design provides the methodology for testing which includes three types of tests that need to be conducted: (1) Competency question verification, (2) Inference verification, and (3) Error provocation tests. Competency question verification consists in the reformulation of the competency questions from natural language to SPARQL queries and running them against the ontology using a toy dataset which is supposed to include the expected result of the query. Inference verification tests are used to check the inferences over the ontologies, by comparing the expected inferences to the actual ones. Lastly, error provocation is a stress test of the ontology to verify how the ontology reacts when it is fed with erroneous facts or boundary data. To formulate and execute the tests, ontology testers, that have not been involved with the modelling of the module, are assigned in pairs for each module. The protocol that is followed for the testing of the ontologies based on XD is presented in figure 4.1.

1	Gather requirements	For a specific type of test, retrieve all the requirements of the current module that are relevant to this test type.
2	Select requirement	Following the principle of unit testing, select one requirement to test in each test case.
3	Formulate test procedure	Determine how to test that particular requirement.
4	Create test case	Create the test case OWL-file, and an additional OWL-file for storing the first test run and describe both using the test case metamodel and its properties.
5	Add test data	Add the test data, needed to perform the procedure according to step 3, in the test run OWL-file.
6	Determine expected results	Depending on the test data, what would be the output of a correct test run?
7	Run test	Execute the procedure from step 3 on the test run OWL-file with its data from step 5, and record the results.
8	Compare result	Verify the expected output (step 6) against the actual result from step 7.
9	Analyze unexpected result	If the result is not the expected one, analyze why and document any change suggestions or issues.
10	Document	Store all information about the test run and its related test case by using the properties of the test metamodel.
11	Iterate	If there are more requirements of this module to test, return to step 2.

Figure 4.1: Testing methodology in XD.

4.1.1 Towards testing automation in XD

In this subsection we present the work progress regarding the testing automation, part of the task 2.5 of this work package. The work consists on developing a GitHub action for the automation of the construction, running and documentation of tests. The first type of test that is being automated is the competency question verification. The automation is based on five main actions which complete different tasks, such as:

1. Environment setup
2. Input crosscheck

3. Test case construction
4. Running of the test
5. Report/Document the test

GitHub actions provide an interesting infrastructure for the automation of the testing, which is currently a manually executed job. There are several challenges that we are facing such as the refinement of the XD testing methodology to accommodate automation, and the involvement of the ontology engineers in providing input during an Action run. Technically speaking, the latter is a major challenge because it defies the purpose of the automation which is to not include users in the process. We need to do so because we need input from the engineers about the tests that need to be executed. Nonetheless, the state of the project, where selected modules are stable for testing, offers the perfect environment to test the automation and collect new requirements for continuous development.

4.1.2 Polifonia music expert community (PMEC)

Despite the fact that the evaluation of the PON will be certainly performed through the pilots and their demonstrators, and that ontology testing [32] (which is different from ontology evaluation [31]) as shown in the previous section can partially serve as evaluation as part as domain coverage is concerned; we have also laid the foundations for the Polifonia music expert community (PMEC). This community has the objective of driving, orienting, and refining the proposal of requirements from the pilots, and a key role as evaluators of the PON. Despite their participation in some analysis of CQs, the role of the PMEC is to strictly act as evaluators of the PON. The PMEC community leader will make sure that all personas have been assigned to musicologists who are experts on the particular topic addressed by the story. During the ontology design activities, each module in the ontology network will be assigned to an ontology designer (a person in our team) and to an expert team for testing (an expert team includes both a music expert and an ontology designer). Periodically, the testing team will stress and criticise the module's work to ensure that concepts and relationships are sound and well-defined. Nevertheless, an ontology designer might struggle with some concepts requiring a strong music background. To address this problem, we will have a mailing list and a Discord channel. In this way, we won't slow down the ontology design activities as the ontology designer's request can be timely sorted out by the experts. In conclusion, the PMEC members have three main areas of responsibility:

1. **Story involvement** Assignment of a story to each expert, based on domain of expertise. Tasks include analysis of competency questions.
2. **Module-based involvement** Assignment of an ontology module to each expert for evaluation. Task focused on criticising the module's work to ensure that concepts and relationships are sound and well-defined.
3. **"On-demand" involvement** Respond to mailing list and Discord channel for asynchronous solving of doubts by the ontology design task force.

The domain expert community has been involved with the creating of the stories and the formulation of the competency question, from the beginning of the project. Considering the experience throughout this first year, we are working towards the involvement of the music experts as proponents of requirements, and not the involvement in the formulation of competency questions. Music experts will provide examples of data that they need to be represented in the Polifonia ontology network and knowledge graph, and from which competency questions can be derived. This task will be solely responsibility of the ontology engineers, who consistently work with the XD methodology. They can ensure that the practices of XD are followed strictly and most importantly that the requirements are precise and in no need for further analysis or clarification from the music experts.

This community is under the development, and we are currently in the process of developing an interaction framework with effective communication channels to coordinate and keep track of the experts' activities. It is indeed crucial to ensure that the background of music experts is line with the subjects/topics of each ontology module, so that their knowledge and involvement are leveraged to the fullest.

4.2 Use case: Sonar

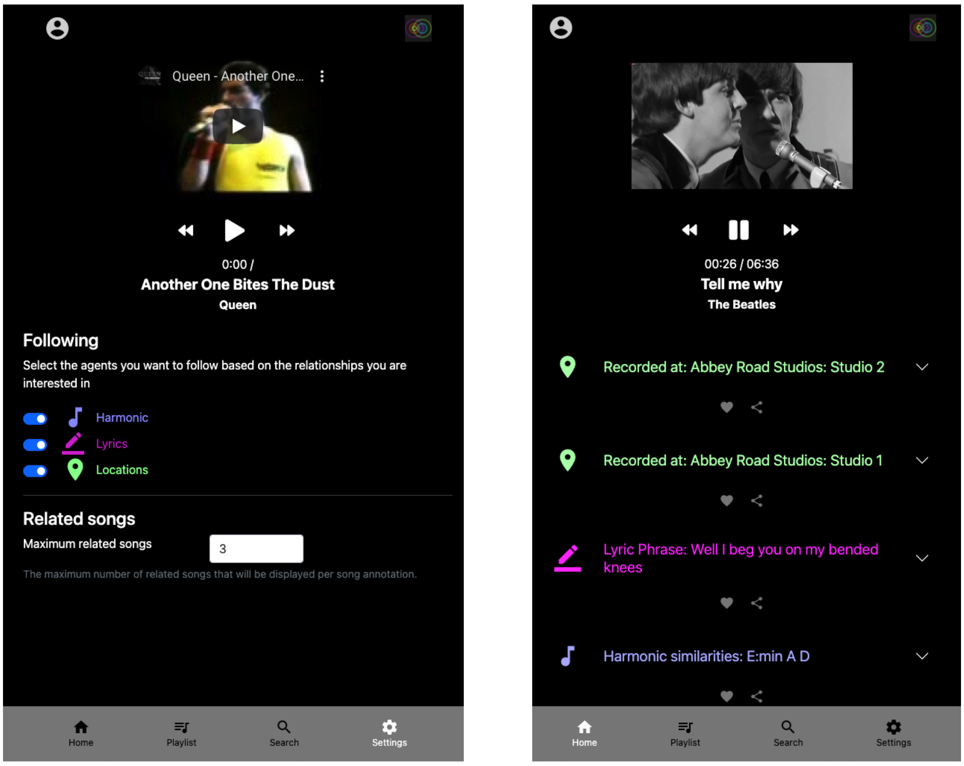

(a) Settings of the Sonar App

(b) The Sonar App displaying information on a Beatles' song

Figure 4.2: The Sonar Demo App

The Polifonia consortium committed to a demonstration of Polifonia technology at the AI and Music festival by SONAR[1] in October 2021 in Barcelona. This presented an excellent opportunity to demonstrate the benefits of knowledge graph creation/integration and knowledge extraction for the musical community.

The Sonia persona[2] profile was the basis for the SONAR experiment, in which a music recommender demo was developed. Three datasets were selected: Isophonics [33], Audio-aligned jazz harmony dataset (JAAH) [34] and Schubert Winterreise [35], to provide a diverse selection of music. Additional facts, relations and metadata were obtained from Wikidata[3], Songfacts, SecondHandSongs[4] and Genius[5]. The music itself was analysed for harmonic and lyric similarity. An app was created that allowed the user to listen to music from the combined datasets. While playing a specific musical piece, links to related pieces of music appeared, based on locations in common, and similarity of harmonics and lyrics. The user could then add these pieces to their playlist to listen to later.

The focus of WP2 was on modelling and testing the ontology modules required to represent these multiple, diverse data sources in a common Knowledge Graph in such a way that related music could be discovered. The Comparative Measure and Musical Performance modules developed for the Ontology Network were extended and refactored for this purpose. Building the demo provided the opportunity to run through the entire eXtreme Design methodology, from CQ based requirements collection through to testing and integration.

The key benefits of the music recommender compared with existing recommenders' functionality were:

- Links were based on objective relations between music as opposed to subjective links, such as what other users also listened to.
- Links were motivated, so the user knew on what basis a piece of music was recommended.
- Links included provenance information, so the user could examine the underlying data used to generate the links.

These benefits are a direct result of the knowledge graph approach, using ontologies to model relations explicitly, transparently and meaningfully, as opposed to a black-box AI approach.

The work was focused around beginning the development of three "intelligent agents" envisioned as artificially intelligent bots. Each of the agents focuses on discovering knowledge of a specific type: one focuses on detecting places and location-related knowledge, another one on detecting similarities between recordings based on the text of their lyrics, and a third one on detecting similarities between recordings based on harmonic similarities by comparing chord progressions in the pieces.

The location agent, provides notes to the user about geographical or social places that are in

[1] https://www.starts.eu/agenda/aimusic-festival-sonar-cccb/detail/
[2] https://github.com/polifonia-project/stories/tree/main/Sonia:%20Playlist%20User
[3] https://www.wikidata.org/wiki/Wikidata:Main_Page
[4] https://secondhandsongs.com/
[5] https://genius.com/

some way connected to the recording playing, such as the place where it was recorded, edited, and mixed at, as well as the birth place of its artist. This data was collected from MusicBrainz[6], Wikidata[7] and SongFacts[8].

The lyrics similarity agent provides notes to the user letting them know about other recordings that have similar lyrics to the recording playing. This data was collected from Genius[9] and SongFacts , and the lyrics similarity was based on state-of-the-art algorithms in the field of Natural Language Processing (NLP). More specifically, to calculate this similarity, all the lyrics of the songs in the dataset were converted into embeddings using LASER[10]. On these embeddings the cosine similarity was then calculated, both at the level of the whole song, and comparing sentence by sentence. .

The harmonic similarity agent provides notes to the user letting them know about other recordings that have similar chord progressions to the recording playing. This data was provided directly from the datasets chosen for this experiment, namely Isophonics [33] for pop-rock music, Audio-aligned jazz harmony dataset (JAAH) [34] for jazz music, and Schubert Winterreise [35] for classical music. , and similarity was calculated using a function specifically designed for this project. In particular, the chords of each song were transposed in the same key, then decomposed into their constituent notes and transformed into numerical values. N-grams were then calculated to identify only those harmonic sub-sequences that were repeated at least once within the song. Then the similarity was calculated by comparing the sub-sequences of all the songs.

The data used by the Sonar App was first of all transformed into RDF, more precisely into Turtle RDF serialisation. To do this, some of the modules of the Polifonia ON were used, namely: the Core Module, the Comparative Measure Module, the Musical Performance Module, the Musical Composition Module, and the Musical Feature Module. All three agents rely on the Musical Performance module to refer to particular musical pieces, more specifically the class `mp:MusicalPerformance` which is in relation with the classes `core:Agent` and `core:Place`, respectively describing the performer of the piece and the place of the performance. In addition, some classes of the Musical Performance Module are used to represent the recording of the performance, such as the classes `mp:Recording` and `mp:Release`, which describe the cataloguing information of a release and the recordings contained therein. The classes `mp:Session` and `mp:SessionType` describe the specific sessions and their types (Recording Session, Mixing Session, etc.).

To represent the harmonic and textual similarity, the Comparative Measure module was designed and developed. For the harmonic similarity, the central class is `cm:HarmonicSimilarity`, which relates it through the object properties `cm:hasEncoding` and `cm:hasSimilarityFunction` to the classes `cm:Encoding` and `cm:SimilarityFuncion`, which describe the encoding and the similarity function used to obtain a certain similarity value, expressed by the datatype property `cm:similarityScore`.

[6] https://musicbrainz.org/
[7] https://www.wikidata.org/
[8] https://www.songfacts.com/
[9] https://genius.com/
[10] https://github.com/yannvgn/laserembeddings

Each harmonic similarity therefore takes into account exactly two entities belonging to the class `cm:ChordProgressionInSimilarity`, which groups the similarity information related to each track. This information consists of the song itself, the harmonic sequence that links the song to another and the temporal information related to the harmonic sequence, expressed by the classes `mp:Recording`, `mf:ChordProgression` and `core:TimeInterval`, respectively.

5 Towards the Polifonia Knowledge Graph

The Polifonia Knowledge Graph uses the ontologies of the PON to annotate and interlink large European music collections as a Knowledge Graph. The results of this task will be reported at a later stage in deliverable *D2.4: Methods for interlinking knowledge graphs – 1st version*. Here, we provide a preliminary summary of ongoing efforts.

Having a comprehensive overview of the landscape of musical data that can be accessed (Section 6.3), together with the ontological ecosystem to represent their content at different levels – from metadata and sub-collections (e.g. album, release), to individual instances of tracks, compositions, artists, and so forth – is a fundamental precondition for the creation of knowledge graphs. According to the scope and the organisation of the project, datasets will also be contributed by the different pilots in Polifonia. For those released as Linked Open Data, using the Polifonia ontology network, whereas datasets provided in different formats (e.g. CSV, JSON, etc.) need to be first semantically described before any integration is possible. Therefore, starting from the pilots, the creation of the Polifonia KG is organised in the two steps outlined as follows.

5.1 From pilot datasets to a Knowledge Graph

If all music collections were provided according to a unique standard, regulating the publication of musical resources on the Web, their integration within the same model and data infrastructure would be simple to implement with current methods – especially if data is described using Semantic Web technologies. However, not all the data collections used within the project Polifonia are provided with the same standard, structure, and format. To overcome this issue, we decided first of all to organise all the data used in the project in a unique directory. Hence, we created a GitHub repository[1] that contains references to all the data collection used. For each of the collection, we create a folder in the directory containing all the information about the dataset, such as the license, the type of data contained, and the instructions for downloading the data. Moreover, each of the dataset folder contains the metadata of the dataset (if available), together with the additional metadata extracted using an external tool such as Cornucopia, a toolkit for enriching and integrating information on Music Information Retrieval (MIR) datasets. The dataset folder will contain different types of data that can be generalised as following: Linked data and other formats. The fist category groups all the resources already encoded triples (and hence in any of the available rdf serialisation). In this case, the data will be aligned to the Polifonia Ontology Network, searching for equivalences and similarities. On the other side, for the data not encoded RDF, we will need to encode them according following the schema proposed by the Polifonia Ontology Network. Technically, the creation of the KG as output of this conversion using

[1] Datasets repository available at: https://github.com/polifonia-project/datasets

the ontological specifications, will be implemented with SPARQL Anything [36], a state of the art system for Semantic Web re-engineering that allows users to query arbitrary files (with SPARQL) and granularly convert the results of these queries to KG elements.

5.2 Enriching and interlinking music datasets

The next step is to extend the Polifonia Knowledge Graph and inter-link pilots' datasets – by leveraging online music catalogues, ad-hoc collections, and music-related resources, and exploiting state of the art methods for entity linbking and computational methods for music similarity. Online music catalogues include MusicBrainz[2], Discogs[3], and Muziekweb[4], to name a few. In the same vein, music-related databases like Genius, SongFacts, and SecondHandSongs also provide heterogeneous sources of information contextually related to artists, composers, tracks, and lyrics. In addition, a number of datasets in music information retrieval (MIR) (c.f. Section 6.3.2) can provide high-quality annotations of musical features – including melodic and rhythmic patterns, chord progressions, musical structures, cadence points, tonalities, etc. – contributed by musicologists and music experts. All this data is particularly appealing considering the scope of Polifonia.

To leverage all this heterogeneous and diverse amount of data, we are developing a modular system that interacts on-request with these collections to find, match and link content related to a given content metadata of a dataset (a list of pieces of tracks it covers) and uses state of the art methods for Semantic Web re-engineering to extend the Polifonia Knowledge Graph.

Following the former step, all the different datasets harmonised within the same KG are then compared with each other. The goal here is to establish links across these (now enriched/extended) datasets according to whether they share common entities or musical properties. In particular, this will be implemented at two levels: alignments between vocabularies will define new links between the various collections; additional links will be established based on their harmonic, melodic, rhythmic and structural similarities, whenever the music is available in either symbolic or audio formats (scores, recordings, etc.). Concerning the second level of interlinking, a first implementation of harmonic similarity – which will be used to link pairs of musical objects sharing non-trivial harmonic properties, has been developed in synergy with WP3 as documented in [37] (Chapter 6).

[2] `https://musicbrainz.org`
[3] `https://www.discogs.com`
[4] `https://www.muziekweb.nl`

6 Related Work

The related work consists of ontology engineering methodologies, existing ontologies that partly encode the domains entailed by the pilots and requirements that we build on [1], and existing datasets conveying such requirements in practical settings from the Polifonia consortium and the Music Information Retrieval (MIR) community.

6.1 Ontology engineering methodologies

Various ontology engineering methodologies have been proposed over the years. More recently, due to the inception of the Web and an increasingly connected world, the focus of ontology engineering methodologies has moved towards collaborative ontology engineering [38]. This section provides a high-level summary of the most important ontology engineering methodologies.

Uschold and King (1995). In [39], Unschold and King propose a set of requirements for a methodology for building ontologies, which include the following stages: 1) identify Purpose; 2) building the ontology (ontology capture, ontology coding); 3) integrating existing ontologies; 4) evaluation; 5) documentation. However, they assert that "it may not be necessary to identify competency questions before building the ontology"[1], thus not providing a supporting tool for collecting the requirements, and no clear guidelines nor methods are defined for the evaluation step.

METHONTOLOGY defines a development process that involves the specification, conceptualization, formalization, implementation, and maintenance steps [40]. It refers to [39] for the requirements elicitation step, needed for conceptualizing the domain, and it does not recommend best practices for the ontology development, besides reusing and integrating existing ontologies.

DILIGENT. This ontology engineering methodology [41] comprises the following main activities, which can involve both domain experts, users, knowledge and ontology engineers: building an initial ontology; local adaptation by the users; analysis and revision of the changes and feedback collected from the users; local update once a new version of the ontology is officially released. This methodology takes into account collaborative aspects and the involvement of different stakeholders. However, it does not provide any guidelines for the design of the ontology, and is not test-oriented.

eXtreme Design. This agile ontology engineering methodology [16, 17] focuses on, and provides guidelines for, the reuse of ontology design patterns (ODPs), i.e. small reusable modelling

[1] A competency question is a question, which translates a modelling issues, and that the ontology needs to be able to answer to.

solutions to recurrent modelling problems. It is iterative and incremental, involves different actors and teams, and is strongly test-based. You can find more details in Section 2.1.

NEON. The NEON ontology engineering methodology does not prescribe a rigid workflow, but instead it suggests a variety of pathways for developing ontologies [42]. To do this, it proposes nine scenarios and further guidelines for reach of them for building ontology networks and general knowledge resources. These scenarios stem from all possible combinations between implementing ontologies from specifications from scratch without reuse, and reusing (or not), re-engineering (or taking as they are), merging (or keeping separate), restructuring, and localizing existing ontological and non-ontological resources and ontology design patterns. On the basis of this, NEON proposes two ontology network life cycle models.

SAMOD. The Simplified Agile Methodology for Ontology Development[2] [43] is an agile methodology that supports the development of an ontology by means of several small and iterative steps involving both domain experts and ontology engineers. The 3 iterative steps are: 1) development of a *modelet* formalising a specific (sub)domain, based on a motivation scenario, with the respective test case; 2) merge of the new modelet with the current model; 3) refactoring of the whole new model. The release of a new milestone at each iteration depends on whether all tests are passed. Additionally, SAMOD defines a set of principles that should guide the ontology design, such as using self-explanatory names for the ontology entities and taking into consideration existing ontologies and ontology design patterns. However, it does not provide explicit guidelines nor support for the requirements elicitation and ontology testing steps.

6.2 Survey on available ontologies

Ontologies play a vital role in the representation and management of knowledge, by providing common vocabularies to describe resources and express requests. In the last two decades several ontologies have been developed for diverse music-related applications, dealing with both symbolic notations and audio signals at different levels of specificity.

Ontologies play a vital role in the representation and management of knowledge, by providing common vocabularies to describe resources and express requests. As outlined in Table 6.2, in the last two decades several ontologies have been developed for diverse music-related applications, dealing with both symbolic notations and audio signals at different levels of specificity. Some ontologies have been designed for describing high-level music-related information, such as the The Music Ontology [7] and the DOREMUS Ontology [8]. One application of Semantic Technologies to the musical field is the ETree project [44], which consists on a linked data set exposing meta-data from the Internet Archive Live Music Archive, containing over 17,000,000 triples describing 100,000 performances by 4,000 artists.

Other ontologies describe musical notation, both from the music score and the symbolic points of view. For example, the MIDI Linked Data Cloud [10] proposes the interconnection of symbolic music descriptions encoded in MIDI format, and the CHARM ontology [45] aims to describe

[2] https://essepuntato.it/samod/

musical structures based on the CHARM specifications. The Music Theory Ontology (MTO) [46] aims to describe theoretical concepts related to a music composition, while The Music Score Ontology (Music OWL) [9] represents similar concepts with a focus on music sheet notation. Finally, the Music Notation Ontology [47] focuses on the core "semantic" information present in a score. Other ontologies aim to describe specific aspects of the musical domain, such as the Chord Ontology, the Tonality Ontology, the Temperament Ontology [48], and the Segment Ontology [49].

Other works describe audio signals or the procedures used to produce them, like The Audio Features Ontology [50], The Studio Ontology [51], and The Audio Effects Ontology [52]. Similarly, the Computational Analysis of the Live Music Archive (CALMA) [53] project aims to link metadata of music tracks with computational analyses of these recordings through feature extraction, clustering and classification.

Additionally, ontologies have been used to model listeners' habits and music tastes, as well as similarities between different musical pieces. The COMUS Ontology aims to represent users' musical preferences and context [54]. Similarly, other ontologies were developed for music recommendation systems, as the Uniemotion: the Emotion Ontology [55] and the Similarity Ontology [56]. Finally, the Mobile audio ontology is a semantic audio framework for the design of novel music consumption experiences on mobile devices [57].

However, these ontologies focus on only some of the elements of musical content. However, music consists of a dense connected network of heterogeneous elements (harmony, melody, lyrics, cultural information, etc.) that concert with each other. In contrast, many of these ontologies were developed as stand-alone projects, with few alignments to existing ontologies of the same domain. Moreover, most ontologies are outdated projects: 47% were developed more than 10 years ago, while 70% are more than 5 years old. Most importantly, many projects and ontologies appear not to be longer maintained nor well documented, with some of the URIs unavailable, which hinders their reuse. The Polifonia project aims to tackle these problems by proposing a network of ontologies that can integrate heterogeneous elements related to the musical content into a modular yet unified architecture.

Ontology	Prefix	Description	Domain	Scope	Last up.	IRI	Reference
Music Ontology	mo	Describing musical artefacts for cataloguing purposes (e.g. artists, releases and tracks).	symbolic	high-level	2013	http://purl.org/ontology/mo/	[7]
Chord Ontology	chord	Provides a common, versatile vocabulary for describing chords and chord sequences following the Harte notation.	symbolic	mid-level	2007	http://purl.org/ontology/chord/	[48]
Tonality Ontology	tonality	Provides high-level and low-level descriptors for tonal content in RDF.	symbolic	mid-level	2008	http://purl.org/ontology/tonality/	[48]
Music Instrument Taxonomies	n/a.	Two interpretations, hence two taxonomies, to categorise musical instruments.	general	mid-level	2011	n/a.	[28]
Similarity Ontology	sim	Concepts and properties for describing similarities between things in the RDF/OWL framework.	general	mid-level	2010	http://purl.org/ontology/similarity/	[56]
Segment Ontology	n/a.	An ontological framework to annotate music to perform musicological analysis of segments in a piece.	symbolic	mid-level	2011	n/a.	[49]
MusicOWL - Music Score Ontology	mso	A comprehensive vocabulary for annotating music scores for melodies, dynamics, and tonalities.	symbolic	low-level	2017	http://linkeddata.uni-muenster.de/ontology/musicscore	[9]
Music Theory Ontology	mto	An extension of music ontologies to include the "missing" theoretical concepts-information.	symbolic	low-level	2018	http://purl.org/ontology/mto/	[46]
Temperament Ontology	tm	Models the main concepts, relationships, and parameters of musical temperament, and facilitates the description and inference of various characteristics of specific temperaments.	symbolic	mid-level	2011	http://purl.org/ontology/temperament/	[58]
Music Notation Ontology	mn	A general ontology for representing notated music, with MEI support.	symbolic	low-level	2017	http://cedric.cnam.fr/isid/ontologies/MusicNote.owl	[47], [59]
COMUS	comus	Emotion state from context and user preference information	symbolic.	high-level	2009	http://ceai.ajou.ac.kr/ontology/0.9/comus.owl	[54]
DOREMUS	mus	An extension of the FRBRoo model for describing music catalogs.	symbolic	high-level	2017	http://data.doremus.org/ontology	[8], [60]
Audio Features Ontology	afo	Common concepts to represent some features of audio signals. Notions of music structures are also present, as well as music theoretic concepts.	audio	mid-level	2016	https://w3id.org/afo/onto/1.1	[50]
Studio ontology	(several)	An extension of the Music Ontology for music production in a recording studio, collecting the following more specific ontologies: audio recording, audio mixing and editing.	audio	mid-level	2011	n/a.	[51]
ETree	etree	Describing performances and related audio tracks	audio	high-level	2014-2017	https://etree.linkedmusic.org/vocab/	[44]
CALMA Ontology	calma	Describing the relation between an audio track and a blob containing features analysis.	audio	low-level	2014-2017	http://calma.linkedmusic.org/calma	[53]
The Audio Effects Ontology	aufx	Describes audio effects in music production workflows. Can be seen as part of the Studio Ontology.	audio	high-level	2013	https://w3id.org/aufx/ontology/1.0#	[52]
CHARM Ontology	charm	Representing hierarchical structures from symbolic music.	symbolic	low-level	2015	n/a.	[45]
UniEmotion: the Emotion Ontology	n/a.	Categorises tags into positive emotional tags, negative emotional tags, and factual tags.	audio	high-level	2013	n/a.	[55]
MIDI Linked Data Cloud	midi	Models MIDI streams at the event level.	symbolic	low-level	2017	http://purl.org/midi-ld	[10]
Mobile Audio Ontology	n/a.	A semantic audio framework for the design of novel music consumption experiences on mobile devices.	audio	mid-level	2014	n/a.	[57]

6.3 Survey on available datasets

To ensure that the ontology design is aligned with the music datasets publicly available on the Web, a series of activities were carried out to: (i) align with the data collections used in Polifonia's pilots; (ii) overview the current datasets in the field of Music Information Retrieval (MIR), considering the high-quality annotations provided by such sources; and (ii) prioritise the music collections contributed by the Polifonia consortium – to ensure that the in-house expert knowledge on these particular repositories is leveraged as a peculiar asset of the project. Each of these activities is described in the following subsections.

6.3.1 Datasets from Polifonia's pilots

The first goal was addressed through a detailed analysis of the dataset survey initially carried out in the WP1 data survey [1], with the aim of verifying whether the musical context set by the pilots was in line with the competency questions used for requirement collection. The outcome of this analysis raised several technical concerns regarding the plausibility of a subset of competency questions in light of the available data and the state of the art in computational music analysis. This allowed to refine some of the competency questions according to the analytical framework described in Section 2.3, thereby achieving a preliminary alignment among WP2 – as technology provider of the ontological ecosystem and for the transformation of data into knowledge graphs, and Polifonia's pilots. Furthermore, to further align with the latest work carried out in the pilots, WP2 will receive a specific data samples taken from the collections these pilots are actually using.

6.3.2 Datasets in Music Information Retrieval

During the last 20 years, the field of MIR has seen the introduction of an unprecedented number of music datasets, enabling researchers to train and evaluate algorithms for several tasks, from chord recognition and beat detection, to source separation and mood detection. Concerning the second line of activities, we conducted an in-depth literature review involving more than 200 datasets that have been extensively used to accommodate a wide variety of MIR tasks. The motivation behind this literature review twofold: first of all it aims at collecting diverse high-quality data and annotations related to the musical content, and secondly, it aims to understand the limitations of the datasets under analysis and how the Polifonia Ontology Network can help to address them.

From a metodological point of view, our survey is organised in such a way as to describe/catalogue music datasets based on their collection metadata – a list of fields that are usually expected to be found, either implicitly (from websites, files, manuscripts, additional material) or explicitly (stored in a single file) from an MIR collection. In this way, taxonomies can be created from our survey based on these fields, which are individually described as follows.

Music media type(s) If the dataset explicitly provides any musical content, this field is used to describe whether such content is either in audio or symbolic format. A collection can also provide tracks of both formats, which is becoming a prominent trend in the literature [35].

Duration The approximate duration of each tracks (in seconds), in case the collection does not provide full-length audio recordings (the release of audio clips/excerpts is common for audio datasets, as the complete audio material may not be shareable due to copyright). For instance, the popular music with emotional annotations (PMEmo) dataset [61] provides audio excerpts of chorus sections (as this is copyrighted material), whereas full-length tracks are directly available in the Jamendo collection [62].

Audio format The audio format of the musical material, if full or partial recordings are provided, e.g. MP3, WAV, FLAC. More choices are plausible depending on the collection.

Symbolic format The symbolic format of the musical material, if a valid digital representation of musical scores, related to a notated composition or a transcribed performance. Common options include, but are not limited to MIDI, MusicXML, MEI, although the former is generally more popular considering the high availability of musical data in this format [63].

Other media In addition, or in alternative, to the musical material, datasets might also provide additional information and artifacts related to the musical objects. For example, audio features (e.g. Mel-frequency cepstral coefficients, chromagrams) are usually provided when recordings cannot be released [64], although collections can also release complementary information such as the rank of tracks on music charts (e.g. Billboard) [65].

Records The number of music pieces covered by the collections (compositions and/or performances), regardless of the availability of the corresponding musical material. This number can vary from small collections of 25 compositions [35] or 50 performances [66], to large scale datasets providing more than a million tracks [67].

Genres Music datasets have a tendency to specialise on music belonging to a narrow selection of genres and styles, to make them more consistent with the kind of analysis enabled by the data and the annotations they provide (e.g. the adherence to musical form is generally more strict in classical music compared to contemporary music). Therefore, it is fundamental to contextualise music datasets to the specific genres their music pertain to.

Year of release The year when the dataset was first released, without considering when the actual data collection activities started. Revisions to the dataset are also recorded in order to keep track of the major editing activities, and also, to have an approximate measure of how actively maintained the collection is.

Collection metadata Whether the dataset provides metadata at the collection level, trivially including all these fields used for our survey. This is needed because collection metadata can also include additional information that is not covered by the survey, e.g., the name of the project investigator, the university that is in charge of storing the data, etc. Surprisingly, most MIR datasets do not provide collection metadata in a standard, unified and consistent manner, hence this process needs to be done manually for each source (as in our survey).

Content metadata Ideally, dataset should provide a specification of their content – a document containing a list of tuples, where each element provides information specific to a single track, e.g., $(title, artist, release, MusicBrainz\ identifier)$. This last information is fundamental to disambiguate among compositions and/or performances, especially if any linking operation allowing the interconnection of different collections (the main goal of INTERLINK – one of Polifonia's pilots) should be operated or considered. Notably, some datasets do not provide this document – where all such information is made explicit for each track, although content metadata may again be implicitly scattered across multiple sources, such as files, websites, and manuscripts [33].

Annotations A list of all the annotations – the actual core of an MIR dataset, as this information is key to enable training, evaluation and testing of computational models for music retrieval and analysis. Annotations are generally contributed by domain experts (musicologists, composition teachers, etc.) when it comes to the detection/attribution of musical features, such as music structure [68], key (tonic and scale) [69], chord progressions [70], at different levels of granularity (hierarchical, flat) and temporal resolutions (global, local). When the annotation task does not require musical expertise, as it involves more subjective and less theoretical interactions with the musical content, annotations are provided by listeners following specific guidelines and frameworks (the annotation methodology); examples include music-induced emotions [71], and listening habits [72].

MIR tasks According to the music annotations provided, a dataset can enable one or more MIR tasks. For instance, the availability of music emotion annotations [73] makes it possible to train and evaluate methods for music emotion recognition [74], whereas the a dataset providing annotations of cadences [75] can be used for pattern extraction and cadence detection [76]. If the source code of these computational methods is provided, their recognition performance/accuracy is often ranked and recorded with respect to each dataset they were tested on (e.g. PapersWithCode[3]). The connection between music datasets and algorithms for music analysis and retrieval is thus a peculiar aspect in MIR – where datasets are more appealing for their computational potential, rather than for the information itself.

Access Music collections can be fully or partially accessible to the public (open), requested for research following a formal procedure for the release (on-demand), or explicitly declared by the authors/curators as unavailable (closed). In the second case, fees, commissions, or cost of licences can be asked by the curators, as done for the RWC dataset [77].

Online Whether the resource can be accessed online or needs to be manually provisioned.

API Whether the database can be accessed through an application programming interface (API), allowing for specific user-defined calls to retrieve musical content of interest. Trivially, this implies that the dataset is online (see the previous field), and is accessible to the user. The automatic access of music datasets is a crucial problem in MIR, as this motivated the development of tools and libraries that can facilitate the process of data acquisition and pre-processing [78, 79].

License/copyright The type of licence and redistribution information, if explicitly provided by the dataset.

References Links to the official website of the dataset or to a web-page describing its content

[3] https://paperswithcode.com

Curator	Music media type	Other media	Records	Genre(s)	Accessibility	Online	Year	License	Link
The name of the institution and/or any reference to the person(s) curating the collection.	audio, symbolic, digitalised scores, etc. (more may applky)	visual, etc.	approximate number of records (tracks for ex.)	classical, pop, etc.	open, closed, on-demand	yes, no	when it was released	the copyright license of the dataset (eg. CC BY)	link to the dataset page, or paper
Albert Meroño	Symbolic	Text (embedded in MIDI)	~500K	Classical, videogames, pop	Open	Yes	2017	CC0 1.0	https://zenodo.org/record/579603#.YPAwPhMzbm1
Johan Oomen	audio							CC	https://musopen.org/music/performer/european-archive/
Johan Oomen	audio						from 2014	CC	https://soundcloud.com/beeldengeluid
Johan Oomen	audio							CC	https://www.europeana.eu/nl/collections/topic/62-music?page
Danny Diamond	Symbolic		~7K	Irish traditional	Currently undefined	Yes	2015	No licence spec	http://port.itma.ie/welcome
Danny Diamond	Symbolic		~18k	Irish traditional	open	Yes	2001 (continu	No licence spec	https://thesession.org/tunes/
Danny Diamond	Symbolic		1,224	Irish traditional	open	Yes	1964-1999 (pr	No licence spec	http://www.capeirish.com/webabc/
Peter van Kranenburg	Audio and Symbolic	Images, midi, lyrics	~20K	Dutch Song Culture	Open	Yes	2014-	CC	http://www.liederenbank.nl/mtc

Figure 6.1: Overview of music collections curated within the Polifonia consortium.

in as much detail as possible. This is also complemented with a link to any academic manuscript formally describing the data collection activities and the annotation process.

Besides the limited availability of audio data, the survey also revealed two central issues: (i) MIR datasets are commonly provided as independent and isolated collections, with little or no alignment at the metadata and annotation level; (ii) even when tracks/compositions are coupled with universal identifiers (e.g. MusicBrainz IDs, ISRC), there is no direct way to access and link heterogeneous music-related data from online databases, such as Wikipedia, Genius[4], and Songfacts[5]. The disconnect among music datasets jeopardises their potential integration, and hence their extension and the combination of annotations of different kinds. Simultaneously, the low level of linkage with other databases discourages multi-modal research in the field, where the availability of heterogeneous music-related data (text, images, locations, etc.) is an essential asset. Therefore, the most common method to link multiple music collections is to implement complex data collection and integration pipelines as done by [80].

In sum, MIR datasets are particularly appealing considering the scope of Polifonia, as these collections provide high-quality annotations of musical features – including melodic and rhythmic patterns, chord progressions, musical structures, cadence points, tonalities and so forth – contributed by musicologists and music experts. Although these collections are primarily used for training and testing methods for computational music analysis, their annotations are rarely reused outside the computational domain despite their intrinsic value. If the integration problem is addressed with an ontological ecosystem that can represent and describe MIR datasets within the same infrastructure, their high-quality annotations would be preserved and their connection to the corresponding musical objects would enable the exploration of rich and diverse music-related data, and the automatic discovery and extraction of knowledge. Finally, the work conducted in this survey will serve not only to catalogue the available datasets, but also to make different resources interoperable, which thanks to the Polifonia Ontology Network can be handled as a unique corpus.

[4] https://www.genius.com
[5] https://www.songfacts.com

6.3.3 Internally curated collections

To conclude, a last survey was conducted internally to identify the music collections[6] curated by partners within the Polifonia consortium. For this particular case, a subset of the fields detailed in the previous subsection (collection metadata) was preliminary selected, and all the contributors in the Polifonia consortium were invited to provide the collection metadata of any music dataset that had personally curated. Seven different music collections curated by four Polifonia partners resulted from the survey (c.f. Figure 6.1). Of these datasets, three provide audio recordings, four are based on symbolic music and another include both symbolic and audio tracks.

In addition to the collections mentioned before, it is also worth to remark that the Polifonia consortium can also count on NEUMA[7], thanks to the direct involvement of the Conservatoire national des arts et métiers (CNAM) and Irémus in the project. NUEMA is a large digital library providing rare corpora of music in MEI format, that can be easily accessed, browsed, and searched by users. The library also includes utilities for the annotation of musical scores, thereby realising an online platform that can be used to contribute new material to extend the ecosystem.

[6] In this context, the term “music collection” is intended in a more general sense, to denote music-related data.
[7] http://neuma.huma-num.fr/

7 Conclusions

In this report we have described in detail the first version of the ontologies and knowledge graphs of music objects developed in Polifonia in the context of WP2 and Task T2.1. Our motivation stems from the difficulty of finding, reusing, integraitng and combining European musical heritage datasets, which are currently scattered on the Web and hardly adhere to the FAIR data principles [6, 2]. The dataset survey work of WP1 [1] brings evidence to this, and provides a first set of datasets contributed by the consortium members that are, indeed, in dire need of greater integration and interoperability. In WP2 we depart from these datasetes, and we combine them with the stories and personas reported by the Polifonia pilots as sources of ontological requirements [1]. We apply the eXtreme Design methodology [13] to gather, curate, analyse and normalise these requirements into a set of well-formed competency questions that we take as basis to identify and provide a first design of 10 modules for the Polifonia Ontology Network (PON). We describe these in detail, as well as the personas, stories and competency questions that they address. We conduct a preliminary evaluation of the PON based on two approaches: an-expert based and semi-automated way of using the PON to answer competency questions, and a use-case driven application of the PON to solve data integration issues in a real-world demonstration that we prepared for the Sónar Barcelona International Music Festival. We leverage this work to propose a first pipeline for knowledge graph construction based on the PON, which uses the classes and properties reported here to semantically uplift various datasets into a first, preliminary version of the Polifonia Knowledge Graph (PKG).

The construction of the PKG sets the immediate future of this work, which will focus on refining the modules of the PON to increase the ontological coverage of the surveyed datasets (both here and by WP1 [1]) and the ability to answer competency questions correctly. Moreover, classes and properties of the PON will be extended to support the construction of the PKG in two ways: by modelling notions of harmonic and lyrics similarity towards deliverable *D2.4: Methods for interlinking knowledge graphs, 1st version* (due on M18), which will report on instance-based matching methods for large-scale music object interlinking; and by modelling patterns that result from the pattern analysis ongoing work in WP3. Moreover, the preliminary evaluation described here will be extended by the work in T2.5 and deliverable *D2.7: Ontology testing and evaluation report, 1st version* (also due on M18), which will set up infrastructure and methods to automatically validate the satisfactory answering of CQs via automated unit testing and decentralised code management (i.e. GitHub and Continuous Integration). In the longer run, we will integrate these outcomes with those of deliverable *D2.3: Ontologies and knowledge graphs of music object context*, which will enrich the music object ontological knowledge discussed here with entities extracted with Natural Language Processing in WP4; and those of deliverable *D2.6: Ontology of licensing, ownership and conditions of use* providing models for rights management and provenance information. All these outcomes will inform further classes and properties that will be engineered, through the processes established in this deliverable, into the modules of the PON.

Bibliography

[1] T. Bottini, V. A. Carriero, J. Carvalho, P. Cathé, F. Ciroku, E. Daga, M. Daquino, A. Davy-Rigaux, C. Guillotel-Nothmann, M. Gurrieri, P. van Kemenade, E. Marzi, A. M. no Peñuelala, P. Mulholland, E. Musumeci, V. Presutti, and A. Scharnhorst, "D1.1 roadmap and pilot requirements 1st version," European Commission, The Polifonia consortium, Tech. Rep., 2021.

[2] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne *et al.*, "The fair guiding principles for scientific data management and stewardship," *Scientific data*, vol. 3, no. 1, pp. 1–9, 2016.

[3] T. Heath and C. Bizer, "Linked data: Evolving the web into a global data space," *Synthesis lectures on the semantic web: theory and technology*, vol. 1, no. 1, pp. 1–136, 2011.

[4] T. R. Gruber, "Toward principles for the design of ontologies used for knowledge sharing?" *International journal of human-computer studies*, vol. 43, no. 5-6, pp. 907–928, 1995.

[5] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier *et al.*, "Knowledge graphs," *ACM Computing Surveys (CSUR)*, vol. 54, no. 4, pp. 1–37, 2021.

[6] M. Daquino, E. Daga, M. d'Aquin, A. Gangemi, S. Holland, R. Laney, A. M. Penuela, and P. Mulholland, "Characterizing the landscape of musical data on the web: State of the art and challenges," 2017.

[7] Y. Raimond, S. Abdallah, M. Sandler, and F. Giasson, "The music ontology," in *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007)*, Vienna, Austria, Sep. 2007.

[8] P. Lisena and R. Troncy, "Doing reusable musical data (DOREMUS)," in *Proceedings of Workshops and Tutorials of the 9th International Conference on Knowledge Capture (K-CAP2017), Austin, Texas, USA, December 4th, 2017*, ser. CEUR Workshop Proceedings, vol. 2065. CEUR-WS.org, 2017.

[9] J. Jones, D. de Siqueira Braga, K. Tertuliano, and T. Kauppinen, "Musicowl: The music score ontology," in *Proceedings of the International Conference on Web Intelligence*, ser. WI '17. New York, NY, USA: Association for Computing Machinery, 2017.

[10] A. Meroño-Peñuela, R. Hoekstra, A. Gangemi, P. Bloem, R. de Valk, B. Stringer, B. Janssen, V. de Boer, A. Allik, S. Schlobach, and K. Page, "The midi linked data cloud," in *The Semantic Web – ISWC 2017*. Cham: Springer International Publishing, 2017.

[11] A. Meroño-Peñuela, R. Pernisch, C. Guéret, and S. Schlobach, "Multi-domain and explainable prediction of changes in web vocabularies," in *Proceedings of the 11th on Knowledge Capture Conference*, 2021, pp. 193–200.

[12] R. Pernisch, D. Dell'Aglio, and A. Bernstein, "Toward measuring the resemblance of embedding models for evolving ontologies," in *Proceedings of the 11th on Knowledge Capture Conference*, 2021, pp. 177–184.

[13] E. Blomqvist, K. Hammar, and V. Presutti, "Engineering ontologies with patterns-the extreme design methodology," *Ontology Engineering with Ontology Design Patterns*, no. 25, 2016.

[14] G. Klyne and J. Carroll, "Resource description framework (rdf): Concepts and abstract syntax," http://www.w3.org/TR/2004/REC-rdf-concepts-20040210/, W3C, Tech. Rep., 2004.

[15] V. A. Carriero, A. Gangemi, M. L. Mancinelli, A. G. Nuzzolese, V. Presutti, and C. Veninata, "Pattern-based design applied to cultural heritage knowledge graphs," *Semantic Web*, vol. 12, no. 2, pp. 313–357, 2021. [Online]. Available: https://doi.org/10.3233/SW-200422

[16] E. Blomqvist, V. Presutti, E. Daga, and A. Gangemi, "Experimenting with extreme design," in *Knowledge Engineering and Management by the Masses. EKAW 2010*, ser. Lecture Notes in Computer Science, P. Cimiano and H. S. Pinto, Eds., vol. 6317. Springer, Berlin, Heidelberg, 2010, pp. 120–134.

[17] E. Blomqvist, K. Hammar, and V. Presutti, "Engineering ontologies with patterns - the extreme design methodology," in *Ontology Engineering with Ontology Design Patterns - Foundations and Applications*, ser. Studies on the Semantic Web, P. Hitzler, A. Gangemi, K. Janowicz, A. Krisnadhi, and V. Presutti, Eds. Amsterdam: IOS Press, 2016, vol. 25.

[18] E. Blomqvist, A. Gangemi, and V. Presutti, "Experiments on pattern-based ontology design," in *K-CAP '09: Proceedings of the fifth international conference on Knowledge capture*, Y. Gil and N. F. Noy, Eds. New York, NY, USA: ACM, 2009.

[19] M. Gruninger and M. S. Fox, "The role of competency questions in enterprise engineering," in *Benchmarking — Theory and Practice. IFIP Advances in Information and Communication Technology*. Springer, Boston, MA, 1994, pp. 83–95.

[20] G. Lodi, L. Asprino, A. G. Nuzzolese, V. Presutti, A. Gangemi, D. R. Recupero, C. Veninata, and A. Orsini, *Semantic Web for Cultural Heritage Valorisation*. Springer, Cham, 2017, pp. 3–37.

[21] V. A. Carriero, A. Gangemi, M. L. Mancinelli, L. Marinucci, A. G. Nuzzolese, V. Presutti, and C. Veninata, "Arco: The italian cultural heritage knowledge graph," in *The Semantic Web - ISWC 2019*, ser. Lecture Notes in Computer Science, C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. F. Cruz, A. Hogan, J. Song, M. Lefrançois, and F. Gandon, Eds., vol. 11779. Springer, Cham, 2019, pp. 36–52.

[22] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, and L. Schneider, "Sweetening ontologies with dolce," in *International Conference on Knowledge Engineering and Knowledge Management*. Springer, 2002, pp. 166–181.

[23] V. A. Carriero *et al.*, "The landscape of ontology reuse approaches," *Applications and Practices in Ontology Design, Extraction, and Reasoning*, vol. 49, p. 21, 2020.

[24] V. Presutti, G. Lodi, A. Nuzzolese, A. Gangemi, S. Peroni, and L. Asprino, "The role of ontology design patterns in linked data projects," in *International Conference on Conceptual Modeling*. Springer, 2016, pp. 113–121.

[25] C. Guillotel-Nothmann, "Knowledge extraction and modelling in the project thesaurus musicarum germanicarum," in *Prague DH Workshops Session III: Editions*, 2020.

[26] C. Guillotel-Nothmann and A.-E. Ceulemans, "Das diatonisch-chromatische system zur zeit des michael praetorius. eine digitale neuerschließung des syntagma musicum (1619) in verbindung mit dem tanzzyklus terpsichore (1612)," in *Musik im Umbruch. Michael Praetorius zum 400.* Todestag, Harrassowitz, 2021.

[27] A.-E. C. Christophe Guillotel-Nothmann, "Praetorius' polyhymnia caduceatrix (1619): Diatonische logik und skalenlogik in mehrchörigen und konzertierenden werken des frühbarocks," in *Zeitschrift der Gesellschaft für Musiktheorie, Gesellschaft für Musiktheorie, 2020, Alte Musik und Musiktheorie zwischen Historiografie und Pädagogik*, vol. 17, no. 2, 2020.

[28] S. Kolozali, M. Barthet, G. Fazekas, and M. B. Sandler, "Knowledge representation issues in musical instrument ontology design." in *ISMIR*, 2011, pp. 465–470.

[29] P. Groth and L. Moreau, "Prov-overview," *W3C Working Group Note*, 2013.

[30] T. Kuhn, A. Meroño-Peñuela, A. Malic, J. H. Poelen, A. H. Hurlbert, E. C. Ortiz, L. I. Furlong, N. Queralt-Rosinach, C. Chichester, J. M. Banda *et al.*, "Nanopublications: a growing resource of provenance-centric scientific linked data," in *2018 IEEE 14th International Conference on e-Science (e-Science)*. IEEE, 2018, pp. 83–92.

[31] A. Gangemi, C. Catenacci, M. Ciaramita, and J. Lehmann, "Modelling ontology evaluation and validation," in *European Semantic Web Conference*. Springer, 2006, pp. 140–154.

[32] E. Blomqvist, A. S. Sepour, and V. Presutti, "Ontology testing-methodology and tool," in *International Conference on Knowledge Engineering and Knowledge Management*. Springer, 2012, pp. 216–226.

[33] M. Mauch, C. Cannam, M. Davies, S. Dixon, C. Harte, S. Kolozali, D. Tidhar, and M. Sandler, "Omras2 metadata project 2009," in *12th International Society for Music Information Retrieval Conference*, ser. ISMIR, 2009.

[34] V. Eremenko, E. Demirel, B. Bozkurt, and X. Serra, "Jaah: Audio-aligned jazz harmony dataset," Jun. 2018. [Online]. Available: https://doi.org/10.5281/zenodo.1290737

[35] C. Weiß, F. Zalkow, V. Arifi-Müller, M. Müller, H. V. Koops, A. Volk, and H. G. Grohganz, "Schubert winterreise dataset: A multimodal scenario for music analysis," *Journal on Computing and Cultural Heritage (JOCCH)*, vol. 14, no. 2, pp. 1–18, 2021.

[36] E. Daga, L. Asprino, P. Mulholland, and A. Gangemi, "Facade-x: An opinionated approach to sparql anything," in *Volume 53: Further with Knowledge Graphs*, M. Alam, P. Groth, V. de Boer, T. Pellegrini, and H. J. Pandit, Eds. IOS Press, August 2021, vol. 53, pp. 58–73. [Online]. Available: http://oro.open.ac.uk/78973/

[37] J. McDermott, D. Diamond, A. S. Khattak, P. van Krannenburg, J. de Berardinis, A. Poltronieri, A. M. no Peñuela, R. Fournier-S'niehotta, and T. Zhu, "D3.2 analysis of music repositories to identify musical patterns," European Commission, The Polifonia consortium, Tech. Rep., 2021.

[38] E. Simperl and M. Luczak-Rösch, "Collaborative ontology engineering: a survey," *The Knowledge Engineering Review*, vol. 29, no. 1, pp. 101–131, 2014.

[39] M. Uschold and M. King, *Towards a methodology for building ontologies*. Citeseer, 1995.

[40] M. Fernández-López, A. Gómez-Pérez, and N. Juristo, "Methontology: from ontological art towards ontological engineering," 1997.

[41] H. S. Pinto, S. Staab, and C. Tempich, "Diligent: Towards a fine-grained methodology for distributed, loosely-controlled and evolving engineering of ontologies," in *ECAI*, vol. 16. Citeseer, 2004, p. 393.

[42] M. C. Suárez-Figueroa, A. Gómez-Pérez, and M. Fernández-López, "The neon methodology for ontology engineering," in *Ontology engineering in a networked world*. Springer, 2012, pp. 9–34.

[43] S. Peroni, "A simplified agile methodology for ontology development," in *OWL: Experiences and Directions–Reasoner Evaluation*. Springer, 2016, pp. 55–69.

[44] S. Bechhofer, K. Page, and D. De Roure, "Hello cleveland! linked data publication of live music archives," in *2013 14th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS)*. IEEE, 2013, pp. 1–4.

[45] N. Harley and G. Wiggins, "An ontology for abstract, hierarchical music representation," in *Demo at the 16th International Society for Music Information Retrieval Conference (ISMIR 2015), Malaga, Spain*, 2015.

[46] S. M. Rashid, D. De Roure, and D. L. McGuinness, "A music theory ontology," in *Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music*, ser. SAAM '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 6–14.

[47] S. S.-s. Cherfi, C. Guillotel, F. Hamdi, P. Rigaux, and N. Travers, "Ontology-based annotation of music scores," in *Proceedings of the Knowledge Capture Conference*, ser. K-CAP 2017. New York, NY, USA: Association for Computing Machinery, 2017.

[48] G. Fazekas, Y. Raimond, K. Jacobson, and M. Sandler, "An overview of semantic web activities in the omras2 project," *Journal of New Music Research*, vol. 39, 12 2010.

[49] B. Fields, K. R. Page, D. D. Roure, and T. Crawford, "The segment ontology: Bridging music-generic and domain-specific," in *Proceedings of the 2011 IEEE International Conference on Multimedia and Expo, ICME 2011, 11-15 July, 2011, Barcelona, Catalonia, Spain*. IEEE Computer Society, 2011.

[50] A. Allik, G. Fazekas, and M. B. Sandler, "An ontology for audio features," in *Proceedings of the 17th International Society for Music Information Retrieval Conference, ISMIR 2016, New York City, United States, August 7-11, 2016*, 2016.

[51] G. Fazekas and M. B. Sandler, "The studio ontology framework," in *Proceedings of the 12th International Society for Music Information Retrieval Conference, ISMIR 2011, Miami, Florida, USA, October 24-28, 2011*. University of Miami, 2011.

[52] T. Wilmering, G. Fazekas, and M. B. Sandler, "The audio effects ontology," in *Proceedings of the 14th International Society for Music Information Retrieval Conference, ISMIR 2013, Curitiba, Brazil, November 4-8, 2013*, 2013.

[53] S. Bechhofer, S. Dixon, G. Fazekas, T. Wilmering, and K. Page, "Computational analysis of the live music archive," in *ISMIR2014, 15th International Society for Music Information Retrieval Conference*, 2014.

[54] S. Rho, S. Song, E. Hwang, and M. Kim, "Comus: Ontological and rule-based reasoning for music recommendation system," in *Advances in Knowledge Discovery and Data Mining*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009.

[55] H. H. Kim, "A semantically enhanced tag-based music recommendation using emotion ontology," in *Intelligent Information and Database Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.

[56] K. Jacobson, Y. Raimond, and M. B. Sandler, "An ecosystem for transparent music similarity in an open world," in *Proceedings of the 10th International Society for Music Information Retrieval Conference, ISMIR 2009, Kobe International Conference Center, Kobe, Japan, October 26-30, 2009*. International Society for Music Information Retrieval, 2009.

[57] F. Thalmann, A. P. Carrillo, G. Fazekas, G. A. Wiggins, and M. Sandler, "The mobile audio ontology: Experiencing dynamic music objects on mobile devices," in *2016 IEEE Tenth International Conference on Semantic Computing (ICSC)*. IEEE, 2016, pp. 47–54.

[58] D. Tidhar, G. Fazekas, M. Mauch, and S. Dixon, "Tempest: Harpsichord temperament estimation in a semantic web environment," *Journal of New Music Research*, vol. 39, no. 4, pp. 327–336, 2010.

[59] S. S.-S. Cherfi, F. Hamdi, P. Rigaux, V. Thion, and N. Travers, "Formalizing quality rules on music notation. an ontology-based approach," in *International Conference on Technologies for Music Notation and Representation-TENOR'17*, 2017.

[60] M. Achichi, P. Lisena, K. Todorov, R. Troncy, and J. Delahousse, "Doremus: A graph of linked musical works," in *International Semantic Web Conference*. Springer, 2018, pp. 3–19.

[61] K. Zhang, H. Zhang, S. Li, C. Yang, and L. Sun, "The pmemo dataset for music emotion recognition," in *Proceedings of the 2018 acm on international conference on multimedia retrieval*, 2018, pp. 135–142.

[62] D. Bogdanov, M. Won, P. Tovstogan, A. Porter, and X. Serra, "The mtg-jamendo dataset for automatic music tagging," 2019.

[63] C. Raffel, *Learning-based methods for comparing sequences, with applications to audio-to-midi alignment and matching*. Columbia University, 2016.

[64] J. B. L. Smith, J. A. Burgoyne, I. Fujinaga, D. De Roure, and J. S. Downie, "Design and creation of a large-scale database of structural annotations." in *ISMIR*, vol. 11. Miami, FL, 2011, pp. 555–560.

[65] J. A. Burgoyne, J. Wild, and I. Fujinaga, "An expert ground truth set for audio chord recognition and music analysis." in *ISMIR*, vol. 11, 2011, pp. 633–638.

[66] A. L. Selway, "Chordify annotator subjectivity dataset," 2021.

[67] T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere, "The million song dataset," 2011.

[68] J. Paulus, “Improving markov model based music piece structure labelling with acoustic information.” in *ISMIR*, 2010, pp. 303–308.

[69] P. Knees, Á. Faraldo Pérez, H. Boyer, R. Vogl, S. Böck, F. Hörschläger, M. Le Goff *et al.*, “Two data sets for tempo estimation and key detection in electronic dance music annotated from user corrections,” in *Proceedings of the 16th International Society for Music Information Retrieval Conference (ISMIR); 2015 Oct 26-30; Málaga, Spain.[Málaga]: International Society for Music Information Retrieval, 2015. p. 364-70.* International Society for Music Information Retrieval (ISMIR), 2015.

[70] Q. Xi, R. M. Bittner, J. Pauwels, X. Ye, and J. P. Bello, “Guitarset: A dataset for guitar transcription.” in *ISMIR*, 2018, pp. 453–460.

[71] A. Aljanaki, Y.-H. Yang, and M. Soleymani, “Developing a benchmark for emotional analysis of music,” *PloS one*, vol. 12, no. 3, p. e0173392, 2017.

[72] M. Schedl, “The lfm-1b dataset for music retrieval and recommendation,” in *Proceedings of the 2016 ACM on international conference on multimedia retrieval*, 2016, pp. 103–110.

[73] J. S. Gómez-Cañón, E. Cano, T. Eerola, P. Herrera, X. Hu, Y.-H. Yang, and E. Gómez, “Music emotion recognition: Toward new, robust standards in personalized and context-sensitive applications,” *IEEE Signal Processing Magazine*, vol. 38, no. 6, pp. 106–114, 2021.

[74] Z. Zhang, E. Coutinho, J. Deng, and B. Schuller, “Cooperative learning and its application to emotion recognition from speech,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 1, pp. 115–126, 2014.

[75] M. Giraud, R. Groult, E. Leguy, and F. Levé, “Computational fugue analysis,” *Computer Music Journal*, vol. 39, no. 2, pp. 77–96, 2015.

[76] W. Zhang and G. Li, “Detection of multiple micro-drones via cadence velocity diagram analysis,” *Electronics Letters*, vol. 54, no. 7, pp. 441–443, 2018.

[77] M. Goto, H. Hashiguchi, T. Nishimura, and R. Oka, “Rwc music database: Music genre database and musical instrument sound database,” 2003.

[78] R. M. Bittner, M. Fuentes, D. Rubinstein, A. Jansson, K. Choi, and T. Kell, “mirdata: Software for reproducible usage of datasets,” in *ISMIR*, 2019.

[79] M. Fuentes, J. Salamon, P. Zinemanas, M. Rocamora, G. Paja, I. R. Román, M. Miron, X. Serra, and J. P. Bello, “Soundata: A python library for reproducible use of audio datasets,” *arXiv preprint arXiv:2109.12690*, 2021.

[80] I. A. P. Santana, F. Pinhelli, J. Donini, L. Catharin, R. B. Mangolin, V. D. Feltrim, M. A. Domingues *et al.*, “Music4all: A new music database and its applications,” in *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*. IEEE, 2020.

Appendix A: Analysis and resolution of CQs

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Paul Mulholland domain experts: N/A
Anna1-CQ1	-	What is the rhythm of the bass line?	design		RHYTHM, BASS LINE	Rythm	
Anna1-CQ2	-	What is the rhythm of the electronic drum kit?	design		RHYTHM, ELECTRONIC DRUM KIT	Rythm	

Figure 7.1: Anna's CQ analysis.

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Eleonora Marzi domain experts: N/A
Carolina1-CQ1	-	Where was a musical composition performed?	design		PERFORMANCE, COMPOSITION, PLACE		-
Carolina1-CQ2	-	In which buildings was a musical composition performed?	design		PERFORMANCE, COMPOSITION, PLACE		-
Carolina1-CQ3	-	Where was a musical composition performed for the first time?	design		PERFORMANCE, COMPOSITION, PLACE		-
Carolina1-CQ4	-	In which buildings was a musical composition performed for the first time?	design		PERFORMANCE, COMPOSITION, PLACE		-
Carolina1-CQ5	-	Which performers (musicians, singers) have performed a musical composition?	design	William1-CQ4 (B)	PERFORMANCE, PERFORMER, COMPOSITION		-
Carolina1-CQ6	-	Which performers (musicians, singers) have performed a musical composition for the first time?	design		PERFORMANCE, PERFORMER, COMPOSITION		-
Carolina1-CQ7	-	In which scores is there evidence of a musical composition?	design		COMPOSITION, SCORE	Evidence	-
Carolina1-CQ8	-	In which historical documents is there evidence of a musical composition?	design		COMPOSITION, HISTORICAL DOCUMENT	Evidence	-
Carolina1-CQ9	-	In which bibliographic references is there evidence of a musical composition?	design		COMPOSITION, BIBLIOGRAPHIC REFERENCE	Evidence	-
Carolina1-CQ10	-	In which historical documents is there evidence of reception of a musical composition?	design		COMPOSITION, HISTORICAL DOCUMENT, RECEPTION	Evidence of reception	-
Carolina1-CQ11	-	In which bibliographic references is there evidence of reception of a musical composition?	design		COMPOSITION, BIBLIOGRAPHIC REFERENCE, RECEPTION	Evidence of reception	-
Carolina1-CQ12	-	Which is the medium of performance of a musical composition?	design		MEDIUM OF PERFORMANCE, COMPOSITION	Medium of performance	-
Carolina1-CQ13	-	Which was the first medium of performance of a musical composition?	design		MEDIUM OF PERFORMANCE, COMPOSITION	Medium of performance	-
Carolina1-CQ14	-	Which instruments are involved in a musical composition?	design	Carolina-CQ12	INSTRUMENT, COMPOSITION		-
Carolina1-CQ15	-	Which voices are involved in a musical composition?	design	Carolina-CQ12	VOICE, COMPOSITION		-
Carolina1-CQ16	-	In which tonality was written a musical composition for the first time?	design		TONALITY, COMPOSITION		-
Carolina1-CQ17	-	Which is/are the composer(s) of a musical composition?	design	William1-CQ4 (A)	COMPOSITION, COMPOSER		-

Figure 7.2: Excerpt of Carolinas's CQ analysis.

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Paul Mulholland domain experts: N/A
David1-CQ1	-	What music was being played?	design		PERFORMANCE, COMPOSITION		x
David1-CQ2	-	Who was listening to the music?	design		PERFORMANCE, COMPOSITION, AUDIENCE		x
David1-CQ3	-	How did the instruments used change over time?	design-visual		INSTRUMENT		x
David1-CQ4	-	What was source of money – e.g. for instruments?	design		INSTRUMENT, SOURCE OF MONEY	source of money	
David1-CQ5	-	Where were the musicians coming from?	design		PERFORMANCE, PERFORMERS, PLACE		x
David1-CQ6	-	Who were teaching the musicians?	design		COMPOSITION, COMPOSER	teaching the musician	
David1-CQ7	-	Where were the places (in which they played)?	design		PERFORMANCE, PLACE		x
David1-CQ8	-	How were instruments sold?	design		INSTRUMENT, SALE		x
David1-CQ9	-	What was the sales pitch?	design		INSTRUMENT, SALE	sales pitch	
David1-CQ10	-	How does the world of brass bands connect with that of elite music?	design		?	connection, elite music, world	
David1-CQ11	-	What did people 'make of', how did they react to, e.g. music?	design	Carolina1-CQ20	RECEPTION, COMPOSITION, PERFORMANCE	clarification on "react"	
David1-CQ12	-	Who was the author of a piece of music? e.g. can we identify it by its incipit?	design		COMPOSITION, COMPOSER	incipit	
	A	Who is the composer of a piece of music?		Carolina1-CQ17	COMPOSER, MUSICAL PIECE		x
	B	Can we infer-detect the composer from the incipit?			COMPOSER, MUSICAL PIECE		
David1-CQ13	-	How is it possible to identify a composer by analysing the formal aspects of a music piece?	design		COMPOSITION, COMPOSER	formal aspects	x
David1-CQ14	-	What was the country of origin of a piece of music?	design		COMPOSITION, PLACE	origin	x
David1-CQ15	-	What is the cause of a change, e.g. increasing number of concerts over time?	design		PERFORMANCE, COMPOSITION	cause of change	

Figure 7.3: David's CQ analysis.

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Paul Warren domain experts: N/A
Keith1-CQ1	-	Can we find coonection (relations) between different artists and music?	design	-	COMPOSITION, ARTIST, INFLUENCE, CONNECTION	connection	
Keith1-CQ2	-	Can we find music which has a non-obvious relation to some other music?	design	-	COMPOSITION, INFLUENCE, CONNECTION	related	
Keith1-CQ3	-	Can we visualize the relations between different media, different music etc?	visual	-	COMPOSITION, MEDIA, GENRE	relations	x

Figure 7.4: Keith's CQ analysis.

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Elena Musumeci domain experts: N/A
Keoma1-CQ1	-	Where is the Bell Tower?	design	Patrizia1-CQ1	BELL TOWER, PLACE		x
Keoma1-CQ2	-	What kind of repertoire is traditionally performed in that specific location?	design	Patrizia1-CQ2	BELL REPERTOIRE, PLACE, PERFORMANCE	kind of repertoire	
	R	What repertoire is traditionally performed in a specific location?			BELL REPERTOIRE, PLACE, PERFORMANCE		
Keoma1-CQ3	-	What is the dating that accompanies a repertoire?	design	-	BELL REPERTOIRE, PLACE, TIME	dating	x
Keoma1-CQ4	-	Is it a repertoire performed in the contemporary world?	design	-	BELL REPERTOIRE, TIME	contemporary world	x
	R	Is it a repertoire currently performed, or was it performed only in the past?			BELL REPERTOIRE, TIME		
Keoma1-CQ5	-	Is it a repertoire in that place currently performed by hand or by electric means?	design	Patrizia1-CQ2	REPERTOIRE, PLACE, SOUND PRACTICE	means	x
Keoma1-CQ6	-	How many people are generally employed for the performance of these sound practices?	design	-	SOUND PRACTICE, MANUAL SOUND PRACTICE	people	x
	R	How many performers (or bell ringers?) are generally employed for the performance of these sound practices?			PEOPLE (PERFORMER), PERFORMANCE, SOUND PRACTICE		
Keoma1-CQ7	-	Are there significant variations in the frequencies between the manually produced sound and the electric sound, such as to have repercussions on the soundscape?	design	-	MANUAL SOUND PRACTICE, ELECTRIC SOUND PRACTICE, SOUND FREQUENCY	repercussion on the soundscape	
Keoma1-CQ8	-	Are there historical evidences of the presence of that sound in the surrounding area?	design	-	HISTORICAL DOCUMENT, BIBLIOGRAPHIC REFERENCE, PLACE	sound	x

Figure 7.5: Keoma's CQ analysis.

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Paul Warren domain experts: N/A
Mark1-CQ1	-	Can we identify a tune (e.g. from an oral tradition) in our collection with music in another documented collection, e.g. RISM, NEUMA, ABC	design	-	TUNE, MUSIC COLLECTION	Tune, identify	
Mark1-CQ2	-	Can we compare music from different collections, e.g. from different countries to show connections/influences between musical styles?	design	-	MUSICAL STYLE, CONNECTION, MUSIC COLLECTION	Connection and influences Music	
	A	Which is the country of origin of a composition?	design	-	COMPOSITION, PLACE		x
	B	Which is the musical style of a composition?	design	-	MUSICAL STYLE, COMPOSITION		x
	C	Which is the collection a composition is member of?	design	-	COMPOSITION, MUSIC COLLECTION		x
Mark1-CQ3	-	Can we compare music longitudinally, e.g. to see evolution of tonality and transition from modal to tonal? Working either at manuscript or collection level.	design	-	TONALITY	Evolution of tonality	
Mark1-CQ4	-	Can we visualize changes in musical form over time, e.g. transition from modal to tonal (evolution curves - see Weiss et al. 2018 'Investigating style evolution of Western classical music: A computational approach') - with zooming	design-visual	-	MUSICAL FORM, TONALITY	Visualise changes	
	R	What is the musical form of a certain composition?		-	MUSICAL FORM, COMPOSITION		x
Mark1-CQ5	-	Can we visualize interconnections, e.g. of tunes which share melodic patterns or geographical origin?	design-visual	-	CONNECTION, MELODIC PATTERN, PLACE	Interconnections, Melodic patterns	
	A	Which tunes share a melodic pattern?	design	-	TUNE, MELODIC PATTERN		x
	B	Which tunes share their geographical origin?	design	-	TUNE, PLACE		x

Figure 7.6: Mark's CQ analysis (first story).

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Marilena Daquino, Enrico Daga, Paul Warren domain experts: N/A
Ortenz1-CQ1	-	What are the aspects that characterize an official source? E.g. context of production, creator's role, relation between the source and political or national events	design	-	SOURCE, CREATOR OF SOURCE, POLITICAL/NATIONAL EVENT	source, context of production, creator's role	
	A	Which is the context of production of an official source?	design		SOURCE, CREATOR OF SOURCE		x
	B	Which is the role of the creator of an official source?	design		SOURCE, CREATOR OF SOURCE		x
	C	Which is the relation between the source and political/national events?	design		SOURCE, POLITICAL/NATIONAL EVENT		x
Ortenz1-CQ2	-	What are the official sources?	design		OFFICIAL SOURCE	official source	x
	A	What are the sources produced by ordinary people?	design		ORDINARY PEOPLE SOURCE		x
Ortenz1-CQ3	-	What are the aspects that characterize the source at hand as trustworthy? E.g. context of production, creator's occupation	design		SOURCE CREDIBILITY		
	A	Which is the occupation of the creator of a source?	design		SOURCE, CREATOR OF SOURCE, OCCUPATION		
	B	Which is the context of production of a source?	design		SOURCE, CREATOR OF SOURCE, CONTEXT OF PRODUCTION		x
Ortenz1-CQ4	-	What are the goals of the source?	design		SOURCE, GOAL	goals	
Ortenz1-CQ5	-	What is the emotional response of music listeners/authors in both types of sources?	design		MUSIC EMOTION	induced vs perceived emotions	
Ortenz1-CQ6	-	In which source there is evidence of children as target audience and music as subject?	design		SOURCE, SUBJECT/GOAL/CONTEXT OF USAGE	evidence	x
Ortenz1-CQ7	-	What are the subjects of the source?	design		SOURCE, SUBJECT		x
Ortenz1-CQ8	-	What are the subjects that are related to national identity?	design		SOURCE, SUBJECT		x
Ortenz1-CQ9	-	What are the subjects that are related to national heritage?	design		SOURCE, SUBJECT		x
Ortenz1-CQ10	-	When was the source produced?	design		SOURCE, TIME		x
Ortenz1-CQ11	-	Where was the source produced?	design		SOURCE, PLACE		x
Ortenz1-CQ12	-	Who produced the source?	design		SOURCE, CREATOR OF SOURCE		x
Ortenz1-CQ13	-	What type of source is it? E.g. Literary text, painting	design		SOURCE, SOURCE TYPE		x
Ortenz1-CQ14	-	What is the context of production of the source? E.g. academic, official source, personal or primary source, commission to an artist	design		SOURCE, CONTEXT OF PRODUCTION		x
Ortenz1-CQ15	-	Are there related events connected to the production of the source? E.g. national fests, war period, political or historical events	design		SOURCE, PRODUCTION, EVENT		x
Ortenz1-CQ16	-	What is context of usage of the source? E.g. propaganda, personal account, review of an event, teaching	design		SOURCE, CONTEXT OF USAGE	context of usage	x
Ortenz1-CQ17	-	How subjects/goals/contexts of usage change over time?	design		SOURCE, SUBJECT/GOAL/CONTEXT OF USAGE, TIME		x
Ortenz1-CQ18	-	What is the goal of the official sources?	design		SOURCE, GOAL	goal	

Figure 7.7: Ortenz's CQ analysis (first story).

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Elena Musumeci domain experts: N/A
Patrizia1-CQ1	-	Where is the Bell Tower?	design	Keoma1-CQ1	BELL TOWER, PLACE		x
Patrizia1-CQ2	-	What kind of execution characterizes that bell tower?	design	Keoma1-CQ2	BELL TOWER, BELL REPERTOIRE	kind of execution	
Patrizia1-CQ3	-	Are there formalized collective actors who carry out sound practices?	design	-	SOUND PRACTICE, COLLECTIVE ACTOR	formalized collective actors, sound practices	
Patrizia1-CQ4	-	Are there formalized transmission methods that characterize the acquisition of skills related to the bell practice?	design	-	BELL REPERTOIRE, TRANSMISSION METHOD	formalized transmission metho	x
Patrizia1-CQ5	-	Which age groups do the transmission practices involve?	design	-	TRANSMISSION PRACTICE, AGE	transmission practices	
Patrizia1-CQ6	-	Do transmission practices involve both men and women?	design	-	TRANSMISSION PRACTICE, GENDER	transmission practices	
Patrizia1-CQ7	-	Do transmission practices include the organization of public events?	design	-	TRANSMISSION PRACTICE, ORGANIZATION, EVENT	transmission practices	
Patrizia1-CQ8	Added	Which is the type of the transmission practice?	design	-	TRANSMISSION PRACTICE, TYPE OF TRANSMISSION PRACTICE	transmission practices	x

Figure 7.8: Patrizia's CQ analysis.

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Christophe Guillotel domain experts: N/A
Sethus1-CQ1	-	To which modes Vide homo has been assigned to in Tonalities? Leonhard, Bernhard, Harold, Siegfried, Frans and Alexander relate the motet to the tonus peregrinus. Robert, however, considers the work to be written in the mixolydian mode.	design		MODE, TONALITY, COMPOSITION		x
	R	Which modes a composition has been assigned to?	design		MODE, COMPOSITION		x
Sethus1-CQ2	-	Are the criteria on which these interpretations are based true in the score and, if so, to what extent?	design		INTERPRETATION, SCORE	true interpretation	x
Sethus1-CQ2.1	-	What are the cadence points in the work?	design		CADENCE		x
Sethus1-CQ2.2	-	What are the part's ranges? Do they correspond to modal octaves?	design		DIATONIC/CHROMATIC REGIONS	part's range	
	A	What are the part's ranges?					
	B	Do the part's ranges correspond to modal octaves?					
Sethus1-CQ2.3	-	What is the final of the bassus?	design		TONALITY		x
Sethus1-CQ2.4	-	What is the last chord?	design		CHORD		x
Sethus1-CQ2.5	-	Are there any melodic patterns related to modality?	design		MELODIC PATTERN, MODE		x
Sethus1-CQ2.6	-	To what diatonic environment does the work belong to?	design		DIATONIC/CHROMATIC REGIONS, TONALITY, SCALE	diatonic environment	
Sethus1-CQ2.7	-	What are the part's clefs and keys?	design		DIATONIC/CHROMATIC REGIONS, TONALITY, SCALE, TONALITIES		x
Sethus1-CQ2.8	-	Are the theoretical models on which these criteria are based – for example the Zarlinian cadential scheme ^1-^3-^5 – in line with a work?	design				
Sethus1-CQ2.9	-	What is the tension between the expected theoretical criteria and their actual realisation in a work?	design			tension	
Sethus1-CQ3	-	Can one assume that the modal ambiguity evidenced by the analytical interpretations is intended?	design			evidenced modal ambiguity	
Sethus1-CQ3.1	-	Where does the work appear within the cycle?	design		COMPOSITION, CYCLE	Lasso's output	
Sethus1-CQ3.2	-	What do we know about this cycle and its place in Lassus' output	design		CYCLE		x
	R	What do we know about this cycle and its place in a composer output?					x
Sethus1-CQ3.3	-	Which poetic text is set to music?	design		INTERTEXTUAL UNITS, ?	poetic text	
Sethus1-CQ3.4	-	Has this poetic text been set to music before?	design		INTERTEXTUAL UNITS, ?	poetic text	
Sethus1-CQ4	-	Apart from Lasso's possible exegetical intentions, what does this work tell us about the status of the modes at this moment in the history of composition?	design		MODE, TONALITY		x
	R	What does a work tell us about the status of the modes at a specific moment in the history of a composition?			MODE, TONALITY		x
Sethus1-CQ4.1	-	Is the modal ambiguity observed here frequent in Lasso's output?	design		MODE	Lasso's output	
	R	Is the modal ambiguity observed frequent in a composer output?					x

Figure 7.9: Excerpt of Sethus's CQ analysis (first story).

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Paul Mulholland domain experts: N/A
Sofia1-CQ1	-	What was the composer's network (patrons, institutions …)?	design	-	COMPOSER	composer's network	
Sofia1-CQ1.1	-	What is the "intellectual framework" of a musician? which explains why certain music was played and how it was played	design		MUSICIAN	intellectual framework	
Sofia1-CQ1.2	-	What is the time relationship between different musicians, e.g. who was working at the same time?	design		MUSICIAN, RELATION, TIME		
	A	At what time was a musician working?	design		MUSICIAN, PERIOD OF ACTIVITY, TIME		x
Sofia1-CQ1.3	-	What relationships exist between institutions, e.g. employment of the same musicians; and between people and institutions, e.g. between a number of composers and a printer?	design		RELATION, INSTITUTION	relationships	
	A	Which is the employer of a musician?	design		MUSICIAN, EMPLOYER		x
Sofia1-CQ2	-	What relationships exist between different sources of information, e.g. payslips and accounts?	design		RELATION, SOURCE	relationships, source of info	
Sofia1-CQ3	-	What parallels are there between the composition of music and the use of language, e.g. in the use of rhetorical delivery?	?		RELATION, COMPOSITION, USAGE OF LANGUAGE	parallels, rhetorical delivery	
Sofia1-CQ3.1	-	What rhetorical strategies are used in a piece of music?	design		COMPOSITION, RHETORICAL STRATEGY	rhetorical strategy	
Sofia1-CQ4	-	Why was a piece of music written to be evocative of something, e.g. the call of a cuckoo?	design		MUSICAL PIECE, EVOCATION	evocative	
Sofia1-CQ4.1	-	What does music evoke?	design	Ortenz1-CQ5?	EVOCATION	evoke	

Figure 7.10: Sophia's CQ analysis (first story).

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Paul Mulholland domain experts: N/A
Sonia1-CQ1	-	What is the musical piece that the persona selects?	design		MUSICAL PIECE		
Sonia1-CQ2	-	With which image/textual information/multimedia/musical piece is the selected musical piece conn	design		MULTIMEDIA, CONNECTION		
Sonia1-CQ3	-	Which is the melodic pattern of the musical piece?	design, visual		MELODIC PATTERN, MUSICAL PIECE		
Sonia1-CQ4	-	In which other musical pieces does the melodic pattern appear?	design		MELODIC PATTERN, MUSICAL PIECE		
Sonia1-CQ5	-	Who is/are the composer(s)/musician(s) of the musical piece?	design		COMPOSER, MUSICIAN, MUSICAL PIECE		
Sonia1-CQ6	-	What information is available for the composer(s)/musician(s)?	design		INFORMATION, COMPOSER, MUSICIAN		
Sonia1-CQ7	-	When/Where was the musical piece recorded?	design		PLACE, TIME, MUSICAL PIECE		
Sonia1-CQ8	-	Has another musical piece "the user likes" been recorded in the same venue?	design		PLACE, MUSICAL PIECE		
Sonia1-CQ9	+	Are there excerpts of the lyrics available for a musical piece? If yes, which are they?	design, visual		LYRIC EXCERPT, MUSICAL PIECE		
Sonia1-CQ10	-	Is the excerpt of the lyric featured in other musical pieces?	design		LYRIC EXCERPT, MUSICAL PIECE		
Sonia1-CQ11	-	Which is the source of the excerpt of the lyrics?	design		SOURCE, LYRIC EXCERPT		
Sonia1-CQ12	+	Is there a person mentioned in the excerpt of the lyrics? If yes, what background information is availa	design		PERSON, LYRIC EXCERPT, INFORMATION		

Figure 7.11: Sonia's CQ analysis (first story).

ID	Breakdown	Competency Question (CQ)	Type	Related CQs	Tags	Ambiguities	Needs approval
Link to the story							creator: Jesse de Vos domain experts: N/A
William1-CQ1		In which collections/datasets does a song occur?	design	-	COMPOSITION, MUSIC COLLECTION		x
William1-CQ2		In which countries does a song occur?	design	-	COMPOSITION, PLACE	occurrence in countries	
William1-CQ3		Which compositions match with William's thematic focus for the exhibition?	design	-		match, thematic focus	x
	A	Which is the thematic focus of a composition?	design	-	COMPOSITION, THEMATIC FOCUS	thematic focus	x
	B	Which is the thematic focus of an exhibition?	design		EXHIBITION, THEMATIC FOCUS	thematic focus	x
William1-CQ4		Which composers and performers are related to these compositions?	design	-	COMPOSER, PERFORMER, COMPOSITION	related	x
	A	Which is/are the composer(s) of a composition?	design	Carolina1-CQ17	See Carolina1-CQ17		x
	B	Which is/are the performer(s) of a composition?	design	Carolina1-CQ5	See Carolina1-CQ5		x
William1-CQ5		In which historical documents is composer X mentioned?	design	Carolina1-CQ8	COMPOSER, HISTORICAL DOCUMENT		x
William1-CQ6		Under which titles is this song known? (multi-language)	design	-	COMPOSITION, TITLE		x
William1-CQ7		Providing a visualisation of relations between the relevant countries, compositions, composers and performers in the various collections and their changes over time	visual	-	RELATION, PLACE, COMPOSITION, COMPOSER, PERFORMER, COLLEC	relations	
William1- CQ8		Provide overviews of the occurrences of search results across certain time periods or genres.	visual	-			
William1- CQ9		What is the relevant metadata for music collections and what do the different fields mean?	design	-	METADATA, MUSIC COLLECTION	relevant metadata	
William1- CQ10		Inspect the degree to which metadata fields have been filled across the collections, see trends and gaps?	visual	-			

Figure 7.12: Williams's CQ analysis.