

The Polifonia Ontology Network: Building a Semantic Backbone for Musical Heritage

Jacopo de Berardinis¹, Valentina Anita Carriero², Nitisha Jain¹,
Nicolas Lazzari², Albert Meroño-Peñuela¹, Andrea Poltronieri², and
Valentina Presutti²

¹ King's College London, 30 Aldwych, London, UK

² University of Bologna, Bologna, Italy
jacopo.deberardinis@kcl.ac.uk

Abstract. In the music domain, several ontologies have been proposed to annotate musical data, in both symbolic and audio form, and generate semantically rich Music Knowledge Graphs. However, current models lack interoperability and are insufficient for representing music history and the cultural heritage context in which it was generated; risking the propagation of recency and cultural biases to downstream applications. In this article, we propose the Polifonia Ontology Network (PON) for music cultural heritage, centred around four modules: Music Meta (metadata), Representation (content), Source (provenance) and Instrument (cultural objects). We design PON with a strong accent on cultural stakeholder requirements and competency questions (CQs), contributing an NLP-based toolkit to support knowledge engineers in generating, validating, and analysing them; and a novel, high-quality CQ dataset produced as a result. We show current and future use of these resources by internal project pilots, early adopters in the music industry, and opportunities for the Semantic Web and Music Information Retrieval communities.

Keywords: Knowledge engineering · Ontology · Music Cultural heritage · Competency questions

1 Introduction

Musical heritage encompasses a diversity of human expressions and experiences leaving heterogeneous traces that are difficult to describe, connect, and preserve [29]. In Europe, music cultural heritage developed through varied sources: musical contents and objects (such as tunes, scores, melodies, notations, recordings, etc.) linked to tangible objects (theatres, conservatoires, churches, instruments, etc.) but also to their cultural and historical contexts, opinions and stories told by people with diverse social and artistic roles (scholars, writers, students, intellectuals, musicians, politicians, journalists, etc), and facts expressed in different styles and perspectives (memoire, reportage, news, biographies, reviews) in different languages (English, Italian, French, Spanish, and German) and across centuries [9]. This diversity creates unique opportunities as well as challenges for researchers and practitioners attempting to study and preserve music heritage.

In the H2020 Polifonia project³, various memory institutions, museums, music archives, scholars, commercial organisations, and citizens ask questions (e.g. “*Which tunes share melodic patterns and geographical origin?*”; “*How do libretto and music relate, e.g. in describing an emotion?*”; “*Can we trace the evolution of tonality and transition from modal to tonal?*”) across these multi-perspective and multi-modal sources. This demands the integration of musicological (notes, chords, modes, theories), historical (events, persons, places, objects) and archival/preservation (metadata, descriptors) data and perspectives. The project comprises 4 cultural institutions (CNAM, NISV, MiC, KNAW) and 10 pilots with a large variety and number of requirements. Ontologies and knowledge graphs (KGs) have the potential to overcome these challenges, and shed light on this wealth of resources by extracting, materialising and linking new music history knowledge that was previously overlooked and therefore missing [44,10].

Various ontologies and knowledge engineering methods have been proposed and applied to music industry and cultural heritage [46,12]. For example, the Music Ontology [46] and DOREMUS [1] provide models to describe music metadata with a focus on discographic and classical music, respectively. Although these ontologies cover some aspects of musical heritage interest, they are individually insufficient to overcome the challenge of integrating the notation, metadata, and historical contexts needed for multi-perspective cultural analyses; thus leaving questions about the relationship between musical theory (melodies, tonalities, chords) and culture (historical events, architecture, geography) unanswered. To date, no available ontological framework integrates music metadata, notation, annotation, source provenance, and cultural heritage object descriptions. To the best of our knowledge, no toolkits exist to support knowledge engineering tasks around the lifecycle of competency questions, which is a central project requirement given the large number of variety of stakeholders, pilots and questions.

In this work, we describe the **Polifonia Ontology Network (PON)**, a **set of new ontologies formalising the semantics of music representation, metadata, annotation, analysis, mediums of performance (instruments), and historical sources (provenance)**, enabling the creation of interoperable knowledge graphs from music datasets. To achieve this, we apply and extend eXtreme Design (XD) [8], a well-known ontology design methodology where ontological requirements are gathered from a comprehensive inventory of competency questions (CQs), and modularity is fostered through the reuse of Ontology Design Patterns (ODPs) [23,31]. We also release the PolifoniaCQ dataset, a collection of 361 competency questions on musical heritage. Further, we validate PON and provide evidence of its current and planned (re)use by three different types of users: (i) the Polifonia pilots, using them to generate musical culture KGs; (ii) a number of industrial and institutional stakeholders and early adopters, planning to use PON to annotate their in-house datasets; and (iii) a survey run in the Semantic Web and Music Technology communities showing intentions of use. More specifically, the contributions of this article are:

³ <https://polifonia-project.eu/>

- *Extensions to XD* centred around CQ extraction and enhancement, including both methodological (a CQ-elicitation framework to mirror use cases from domain experts) and technological (a toolkit for assisted design and iterative improvement of CQs through language models) aspects (Section 3).
- *PolifoniaCQ*, a new dataset of competency questions driving the design and the evaluation of PON, with associated stories and personas (Section 3.1).
- The *Polifonia Ontology Network (PON v1.0)* resources, available on GitHub⁴ and including 15 (CC-BY 4.0) ontology modules (Section 4).
- *Evidence of reuse* and impact from music stakeholders, applications within Polifonia, and interest from various research communities (Section 5).

We also provide SPARQL endpoints of our KGs, and documentation and software to facilitate the understanding, reuse, and extension of PON resources.

2 Related work

Ontologies play a fundamental role in the representation and management of knowledge, by providing common vocabularies to describe resources and queries. In the cultural heritage domain, there have been several efforts in this direction, such as the ArCo ontology, which pertains to the Italian cultural heritage [12], as well as others [10,33]. Several ontologies exist in the music domain for addressing diverse music-related applications, dealing with both symbolic notations and audio signals at different levels of specificity. MusoW⁵ is a catalogue that indexes online music resources, including ontologies and knowledge graphs. Here, we focus on music ontologies and categorise them according to their reference domain: (i) metadata; (ii) music theory; (iii) music notation; and (iv) audio features.

Multiple ontologies exist for describing high-level music-related metadata, with the Music Ontology [46] and DOREMUS Ontology [40] being the most renowned. The MUSO Ontology [56] was proposed to model musical concepts of the Finnish National Library. Other models focus on specific metadata: the Performed Music Ontology⁶ specialises on performances, while the OnVIE Ontology [52] describes mediums of performances. Similarly, the Musical Instrument Taxonomies [38] model instruments conceptually and terminologically, while the Smart Music Instrument Ontology [55] covers sensors and components of instruments within the realm of the Internet of Musical Things [54].

Some ontologies describe different elements ascribable to *music theory*. The Music Theory Ontology (MTO) [47] covers theoretical concepts of music theory, while the Functional Harmony Ontology analyses harmonic sequences through reasoning. The Chord Ontology, Tonality Ontology, and Temperament Ontology [21] cover chords, tonal content, and instrument tuning, respectively.

Ontologies have also attempted to describe *musical notation*. For instance, the MIDI Linked Data Cloud [43] proposes a way to connect symbolic music

⁴ <https://github.com/polifonia-project/ontology-network>

⁵ <https://projects.dharc.unibo.it/musow/>

⁶ <https://performedmusicontology.org/>

descriptions that are encoded in the MIDI format. Meanwhile, the CHARM ontology [28] is focused on representing musical structures. The Music Theory Ontology (MTO) [47] aims to capture the theoretical concepts related to music compositions, while the Music Score Ontology (Music OWL) [35] and the Music Annotation Ontology [16] represent the content of a music score.

Other works focus on audio signals or the procedures used to produce them. For example, The Audio Features Ontology [2], The Studio Ontology [20], and The Audio Effects Ontology [57] are dedicated to describing different aspects of audio production. The Computational Analysis of the Live Music Archive (CALMA) [3] project aims to link metadata of music tracks with computational analyses of 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 [49,37,34,53].

Despite the numerous contributions, the scope of these ontologies is often too specific or ingrained in a genre, style, historical period – often addressing individual music stakeholders and/or datasets. Several ontologies were also developed independently, with little coordination across relevant contributions. In turn, this often hampers reuse and extension, while jeopardising interoperability – an essential requirement for the integration of music datasets [11].

3 The eXtreme Design methodology in Polifonia

The Polifonia Ontology Network (PON) addresses the aforementioned challenges by integrating heterogeneous requirements related to musical content and contexts into a modular yet unified architecture. To develop PON, we rely on, and extend, the eXtreme Design [8,7] (XD) ontology engineering methodology. XD fosters the reuse of ontology design patterns (ODPs) [23,31] and provides support to incrementally address small sets of requirements formalised as competency questions (CQs). This minimises the impact of changes in future releases, which is beneficial to Polifonia (heterogeneous project requirements and participants). Moreover, XD has been successfully applied to the cultural heritage domain [13], and our ontology designers have relevant experience in using this methodology.

The application of XD iterates over a series of steps, for which we detail their process while highlighting our main extensions (see Figure 1).

3.1 Requirements collection

Ontological requirements are collected from *customers* in the form of *user stories* (e.g. “*Tosca* was performed in Rome on 14 January 2000”), which are then translated as competency questions (CQs) – the natural language counterpart of structured queries that the resulting KG should answer [25]. For instance, the previous story may become “*Where was a musical piece performed?*”.

We borrow techniques from User eXperience design [24] to extend this framework with 3 new sections in the story template: persona, goal, and scenario. The *persona* is a research-based description of a typical user: name, age, occupation,

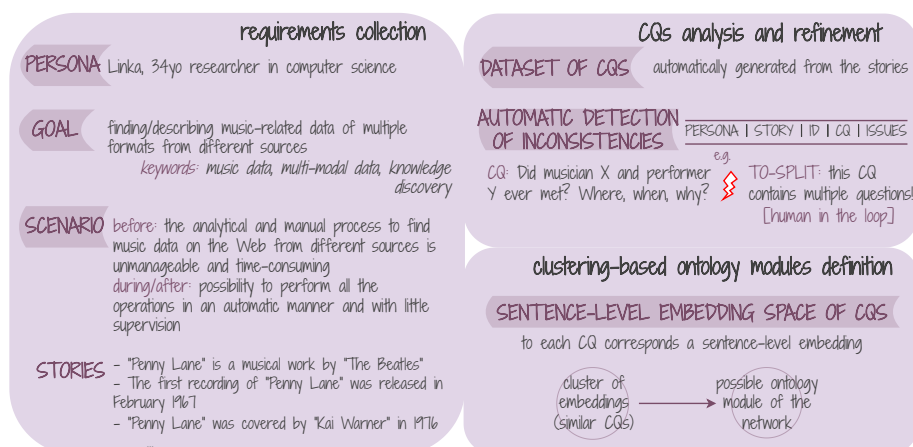


Fig. 1. Summary of the main Polifonia extensions to the eXtreme Design methodology.

skills and interests. The *goal* is a short textual description of what the persona aims to achieve in the story, complemented by a list of keywords (maximum 5) provided by the customers. The *scenario* describes how the persona’s goals are currently solved, to contextualise the gap with the resource being developed.

In cooperation with the domain experts in Polifonia (music historians, librarians, computational musicologists, music analysts, archivists, data engineers, etc.), 22 personas have been created⁷ from this step.

Iterative refinement of CQs. Competency questions were then analysed to identify any inconsistencies that could create obstacles for ontology design. Common inconsistencies were due to vague concepts, for instance, the assertion of two compositions being *connected* without any specific context (in terms of the property) on which the connection can be established (e.g. similar melodies, rhythm). Other CQs were found to be overly complex or nested – entailing more than a single requirement as a result of nested logical operators articulating the question (e.g. “*How is track B connected to C to conclude D?*”). Such CQs needed to be conceptually simplified before being processed further.

To efficiently address these inconsistencies, we developed the *Infer, DEsign, Create* (IDEA) framework: analytical tools for CQ-driven ontology design based on language models⁸. IDEA automatically extrapolates and organises CQs from a source repository, analyses them to find inconsistencies and similarities, and visually projects them to a sentence-level embedding space [51]. The framework has enabled the iterative refinement and improvement of CQs through human-machine collaboration: questions are first extracted and preliminary validated by tagging them (complex, nested, ill-formed, passing), then brought to the attention of the corresponding ontology designer whenever their intervention is

⁷ <https://github.com/polifonia-project/stories>

⁸ <https://github.com/polifonia-project/idea>




Fig. 2. Visualisation of the Polifonia CQ embeddings using TensorBoard.

needed. To date, 3 cycles of CQ improvement have been completed with IDEA. Instead, the analysis of CQ embeddings through similarity facilitated the identification of overlapping requirements from the pilots (beyond the syntactic level); which in turn enabled and fuelled discussion from various experts and pilots during our ontology design meetings (e.g. 2 CQs may have similar interpretation or semantics for OD, but entail different semantics across pilots).

The PolifoniaCQ dataset. At the end of this process, we obtained 361 CQs, which are systematically collected in the PolifoniaCQ dataset with pointers to their personas and stories. We make this dataset available under CC-BY 4.0 ⁹.

3.2 Ontology network design and development

Clustering CQs as ontology modules. The refined CQs could then be translated in clear, atomic and consistent ontological requirements. Given the wide diversity of CQs – ranging from general events to musicological interpretations of specific passages in compositions, the first step was to achieve a meaningful categorisation into thematic clusters. This step led to the definition of the ontology modules shaping the architecture of the Polifonia Ontology Network.

To streamline this process, we analysed the CQ embedding space generated and projected by IDEA. This is done by computing the sentence-level embeddings (a feature vector of fixed size) for each CQ in the PolifoniaCQ dataset. The latter can be considered as a point in a high dimensional space – providing a numerical summary of the question’s meaning [17]. Embeddings are computed via Sentence-BERT [48] due to its state of the art performance on a number of question-related tasks, including multi-lingual search and paraphrase detection.

⁹ <https://github.com/polifonia-project/polifoniacq-dataset>

An interactive visualisation of the PolifoniaCQ embeddings is available from a live Tensorboard Projector [51] which is set up and synchronised via IDEA¹⁰. The qualitative analysis of the embedding space, in addition to density-based clustering analysis under various parametrisations, have jointly facilitated the identification of common requirements (as nested clusters) and enabled the interactive exploration of the PolifoniaCQ dataset via similarity (c.f. Figure 2).

Matching CQs to ODPs. For each module/ontology, an XD iteration starts from selecting a coherent set of CQs. To address those requirements, existing solutions (ODPs) from ontologies or online catalogues of patterns¹¹ are considered for reuse, extension, and specialisation. For instance, a CQ such as “Where and when a situation took place?” can be matched to the *TimeIndexedSituation*¹² ODP, which represents temporal situations.

Here, IDEA supports the identification of “the CQ set” via the multi-lingual search feature. For example, an ontology designer looking for CQs related to places may express a search query as shown below in Listing 1.1.

Listing 1.1. Search results for query “*questions related to places*” with similarity score.

| | | |
|----|-------|--|
| 1 | 0.377 | Where were the places in which musicians played? |
| 2 | 0.368 | Which are all organs near to geographic coordinates x, y? |
| 3 | 0.341 | What are geographically distinct features of organs from a region? |
| 4 | 0.287 | Where is the church/bell tower? |
| 5 | 0.285 | What is the provenance of the event attendees? |
| 6 | 0.275 | Which tunes which share melodic patterns or geographical origin? |
| 7 | 0.265 | What places did a musician visited in her career? |
| 8 | 0.263 | Where is the Bell Tower? |
| 9 | 0.246 | Where was a musical composition performed? |
| 10 | 0.238 | In which buildings was a musical composition performed? |

Direct/indirect ontology reuse. Depending on the project’s requirements, reuse of ontologies and ODPs is direct and/or indirect [14]. The former approach directly includes/imports ontologies or part of them (e.g. individual entities, relations) thus introducing a dependency to any possible changes and availability. In indirect reuse, relevant entities and patterns from other ontologies are used as templates (replicated and extended) while being aligned to ensure interoperability. In Polifonia, we follow a hybrid approach: ArCo ontology [12] is directly reused since its development and maintenance involves one of the project’s partners (MiC), while others (such as DOREMUS) are indirectly reused and aligned.

Validation and testing. Ontology modules have been developed in close collaboration with domain experts and pilot leaders throughout the whole development cycle. This has allowed the ontology design team to leverage the domain expertise in Polifonia to technically validate our modules at different stages: from the collection and analysis of requirements, to iterations of ontology designs. Validation was facilitated by IDEA (at the CQ-level), and, at the modelling level, by the Graphical Framework For OWL Ontologies (Grafoo) notation [19] – providing a powerful visual language for coproduction activities. This has also been

¹⁰ From <https://polifonia-project.github.io/idea/competency-questions/>

¹¹ See for example ontologydesignpatterns.org

¹² <http://ontologydesignpatterns.org/wiki/Submissions:TimeIndexedSituation>

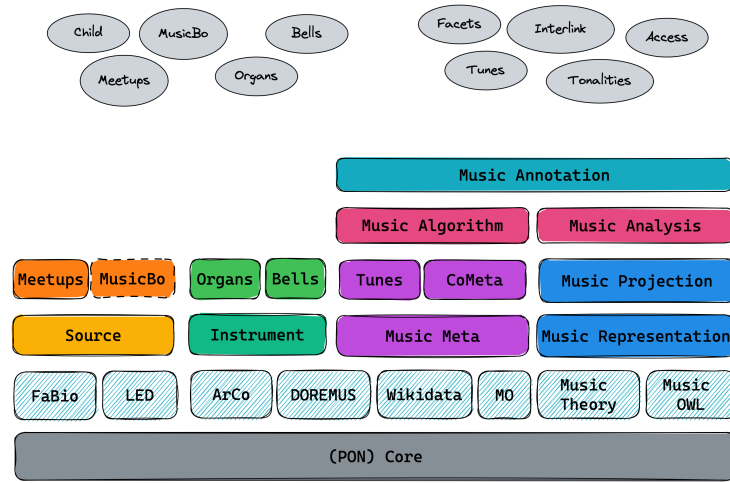


Fig. 3. Overview of the main modules in the Polifonia Ontology Network, with Polifonia’s pilots as early adopters (grey circles). Foundational models (**Source**, **Instrument**, **Music Meta**, **Music Representation**) provide the backbone of PON, built on top of the **Core** module while leveraging the main ontologies reused directly or indirectly.

achieved through data snippets provided by the pilots, which have been modelled by our ontologies and triggered further iterations of improvements.

4 The Polifonia Ontology Network

The Polifonia Ontology Network (PON) provides a modular backbone of music ontologies to address both cultural heritage and more general queries in the music domain. As illustrated in Figure 3, PON v1.0 comprises 15 ontology modules that are organised thematically (colours, horizontal view) and hierarchically, to highlight their dependencies (vertical view). At the bottom of the architecture lies our **Core** module (providing general-purpose elements of design, ODPs, and alignments) and the reused ontologies. Four foundational models provide interoperability across PON through their abstract design: **Source**, **Instrument**, **Music Meta**, and **Music Representation**. These are specialised and extended in the upper levels to add functionalities and contextualise specific domains.

A summary of PON modules is given in Table 1, with links to the repositories storing the modules with documentation, diagrams, and examples. Through our foundational models, PON ontologies can be applied to a wide set of music projects, and the modular design simplifies extensibility and maintenance.

4.1 Foundational models and their extensions and specialisations

The Music Meta module provides a rich and flexible ontology to describe music metadata related to artists, compositions, performances, recordings, broad-

de Berardinis, J. et al. (2023).
The Polifonia Ontology Network: Building a Semantic Backbone for Musical Heritage.
 10.5281/zenodo.7919970

| Module | Prefix | Outline | Repository |
|----------------------|---------|---|--------------------------------|
| Core | core: | Elements of general reuse and ontology design patterns | /core-ontology |
| Music Meta | mm: | Achieving interoperability of music metadata | /music-meta-ontology |
| Music Representation | mr: | Foundational model to describe arbitrary musical content | /music-representation-ontology |
| Music Instrument | mop: | Instruments and their evolution through time and space | /music-instrument-ontology |
| Source | src: | Musical sources and their context of production | /source-ontology |
| Tunes | tunes: | A specialisation of Music Meta for folk music | /tunes-ontology |
| CoMeta | com: | An extension of Music Meta to represent music corpora | /cometa-ontology |
| Music Projection | mp: | Achieving interoperability of music notation systems | /music-projection-ontology |
| Organs | organs: | A rich descriptive model of organs and building methods | /organs-ontology |
| Bells | bell: | Describing bells, bell towers and bell ringers | /bell-ontology |
| Music Algorithm | mx: | Computational methods for music and their parametrisation | /music-algorithm-ontology |
| Music Analysis | ma: | Music analysis through reasoning using modal-tonal theories | /music-analysis-ontology |
| Music Annotation | ann: | A wrapper of ontologies for music annotations (audio, symbolic) | /music-annotation-ontology |
| PON (full) | pon: | The whole Polifonia Ontology Network (imports all modules). | /ontology-network |

Table 1. Overview of the modules in the Polifonia Ontology Network. All URLs are also accessible from <https://github.com/polifonia-project/ontology-network>.

casts, and links. Music Meta focuses on provenance and interoperability – essential requirements for the integration of music datasets, which is currently hampered by the specificity of existent ontologies. The model is based on the Information-Realisation ODP [22], allowing to reduce complexity of FRBR-based models, whose application in the music domain has raised concerns [50].

To enable data integration from existing knowledge bases and datasets, we also align Meta to Music Ontology [46], DOREMUS [2], and Wikidata. To facilitate the reuse of Music Meta and its data conversion into OWL/RDF Knowledge Graphs, we developed a library to map arbitrary music metadata into RDF triples. This enables a practical and scalable workflow for data lifting to create Music Knowledge Graphs without expert knowledge of our ontological model.

The Tunes module extends and specialises Music Meta for folk music. The main novelty consists in grouping and describing *tunes* into “*tune families*” depending on their melodic similarity (an association requiring rich provenance description of the musicological analysis on the source); which also extends to lyrics families.

CoMeta reuses and extends Music Meta to describe arbitrary music collections, corpora, and datasets. Here, metadata is described at the *collection-level* (data curator, annotations provided, availability of audio music, etc.), and at the *content-level*, (e.g., the title, artist, release of each piece in a dataset). The design of CoMeta is informed by a survey of Music Information Retrieval datasets [36].

The Music Representation module provides a comprehensive schema to describe the analysis of musical objects (a score, an audio track, etc.) interpreted according to a theory. Fragments of a musical object (elements of a music object whose temporal location is uniquely identifiable) are described by annotations provided by an agent (e.g. expert annotator, algorithm). An annotation

is either the subjective result of an analysis (e.g. a chord played in a specific section) or objective in nature (e.g. a note in a digital score). Each annotation describes some music content (e.g. notes, chords, etc.), which we refer to as a *musical projection* [41]. Annotations are formalised via our Music Annotation Pattern [4], whereas the definition of music projections is delegated to the **Music Projection** module. The generality of the module and its abstraction over the represented content enables the interoperability of different music annotation schemas. The module is aligned to MusicOWL [35], Music Notation Ontology [15], Music Note Ontology [45], and our JAMS ontology (c.f. Section 4.2).

The Music Projection module formalises musical entities that can be subject of an annotation. This ranges from traditional musical notation (e.g. note, chords) to informal annotations (e.g. mood, danceability). The module is aligned with MusicOWL, Music Notation Ontology, Music Note Ontology, Music Theory Ontology [47], Chord Ontology [21], and Roman Chord Ontology¹³. This allows to integrate existing domain ontologies. Notably, we also harmonise different chord representations (Chord Ontology, the Roman Chord Ontology and the Tonality Ontology) based on the Unified Model of Chords in Western Harmony [30].

The Instrument Module describes musical instruments as mediums of performance and their technical properties. Given that numerous taxonomies of instruments into *groups* and *families* exist (e.g. Hornbostel-Sachs, MIMO, MusicBrainz) and finding common categorisations is an open problem [38], our module provides an abstraction capable to express arbitrary classifications. This is achieved by leveraging the Information-Realisation and the Collection ODPs. Overall, the module allows to: (i) refer to instruments as entities (an instrumentation of a piece for “piano” and “viola”) as well as conceptually (e.g. a viola has 4 strings); (ii) support the integration with different taxonomies and vocabularies, such as [39]; (iii) describe the evolution of instruments in time and space (e.g. a viola as a cultural heritage object being relocated). This provides a foundational level where contributors can “plug” their instrument-specific ontologies [58].

The Bells module extend **Instrument** to describe bells by means of measurable, intrinsic aspects such as weight, materials, conservation status. The main entities contextualising bells are: (i) the author(s), such as the foundry who built the bell; (ii) the agencies that played some role e.g. the agency that took care of cataloguing the bell; (iii) the place(s) where it has been located; (iv) the tower(s) where the bell has been included; (v) the tools that the set of bells is played with; (vi) documents related to the bells, e.g. bibliographies, protective measures.

The Organs module extends **Instrument** to describe organs as (i) a musical instrument consisting of parts; and (ii) as a focal point of *projects* detailing its changes throughout time. To address the former, we used the Parthood pattern

¹³ <https://github.com/polifonia-project/roman-chord-ontology>

from the DOLCE ontology ¹⁴. The entities of the ODP, **Whole** and **Part** make possible the specification of the whole instrument and its parts. In the ontology, the **Whole** entity refers to the organ instrument, and the **Part** entity refers to the parts of the organ that are **Console**, **WindSystem**, **Case**, **Division**, and **Action**.

The Source module represents various sources of music-related information. These include manuscripts, textbooks, articles, interviews, reviews, comments, memoirs, etc. of different scope and format (physical, digital). The module aims to provide general support to describe information related to the *creator* and *type* of the source, the *time* and *place* when/where it was created, the *context of production* and *usage*, and the *subject* and *goals*. Although this conceptualisation leans towards bibliographical sources, the module provides expressivity to indicate multimedia documents (e.g. images of scores, audio recording, video). For example, a video recording of a performance can be considered as a musical source – providing documentary evidence of a composition e.g. during an event.

The Meetups module describes encounters between people in the musical world in Europe from c. 1800 to c. 1945. Historical meetups, which are the main subject of this module, are described by means of four main components: the people involved in the meetup, for instance, the person that is the subject of interest and the people interacting in the event, the place where the encounter took place (e.g., city, country, venue), the type of event, the reason (e.g., music making, personal life, business, among others) and the date when it took place.

The MusicBO module is developed by following a *KG-to-ontology* process [42]. Ontological axioms, grouped into *patterns*, are empirically generated from the MusicBO knowledge graph – which is built from a textual corpus on music performances and encounters between music-related agents in Bologna since the 17th century. Such patterns include information about the probability of axioms to *happen* (as they are derived from the data). For instance, the probability of instances of the pattern *compose* situation (the process of creating art) to have **NaturalPerson** as range of the **artist** property, is higher than the probability of having an **Organisation** as a composer. In sum, the content of the ontology module is highly dependent on the KG, and the most populated and described entities are: persons, places, organisations, works of art, theatres, and books.

4.2 Modules for analysis and annotation of music

The Music Algorithm module formalises algorithms that can operate on music metadata (using the Meta module), and musical content (via the representation module). The module commitments are similar to those defined by Diamantini et al. in [18]. Indeed, an algorithm is characterised by three main components: a *formalisation*, which can be theoretical (e.g. pseudocode) or executable (e.g. using a programming language); a *parametrisation* (e.g. input data); and the kind

¹⁴ <http://www.ontologydesignpatterns.org/ont/dul/DUL.owl>

of *task* it solves. The latter defines a set of entities that are processed alongside the input and output data requirements and the final goal achieved. The module allows theoretical and quantitative performances to be represented in the context of the algorithm’s parametrisation. Through an abstract and general definition, the formalisation in Music Algorithm can be seen as a general pattern, capable of representing any algorithm regardless of the domain of application. In the context of music, the output of the algorithm is considered an analysis, which is then represented via the **Representation** module.

The Music Analysis module allows for the analysis of musical pieces using historical and present-day established musical theories: the *modal* and *tonal* theories. Through the use of this framework, different subjective analyses can be unified – overcoming the limitations imposed by a “global” theoretical perspective. Different theoretical viewpoints can be used for the interpretation of the same piece. Currently, two historical theories are implemented: Zarlino (1558) and Praetorius (1619) [26,27]. Through the use of formal reasoning and a comprehensive axiomatisation, the ontology is able to automatically infer the theoretical interpretations of a musical piece and its evolution in time and space.

The Music Annotation module provides different music annotation models to accommodate musicological and information retrieval use cases. The primary objective of this module is to enhance support for other descriptive systems, thus increasing interoperability and conversion possibilities from various music annotation formats. Indeed, all our models are logically interconnected through **Music Representation**. A fully fledged annotation model here is the JAMS Ontology [6]¹⁵. This ontology mimics the structure of a JAMS (JSON Annotated Music Specification for Reproducible MIR Research) document [32]. It semantically describes and connects all the elements of the JAMS specification (**Annotation**, **Observation**, etc.), including the music metadata and the annotation contents using the **Music Meta** and **Representation** modules, respectively.

5 Adoption and impact

We provide evidence of PON use by Polifonia pilots (Interlink, Tonalities, Meetsups, Bells, and MusicBO), which have contributed 6 musical heritage KGs (Section 5.1); potential interest of reuse and opportunities for the Semantic Web and Music Technology communities collected from an online survey (Section 5.2); early adopters and ongoing synergies from the Polifonia Stakeholder Network for PON validation and annotation of cultural and industrial datasets (Section 5.3).

5.1 Current use by Polifonia pilots

Interlink has published ChoCo and Harmony KGs. Choco [6] provides 20K+ harmonic annotations of scores and tracks, that were integrated from 18 chord

¹⁵ <https://github.com/polifonia-project/jams-ontology>

datasets¹⁶. The KG uses the JAMS ontology in **Music Annotation**, and the **Roman** ontology from the **Music Projection** module. **Harmory** [5] is a KG of interconnected harmonic patterns derived from ChoCo¹⁷, and aimed at human-machine creativity (pattern discovery, chord generation, harmonic similarity). **Tonalities KG** includes data¹⁸ from 377 MEI scores and their annotations w.r.t. theoretical concepts (roots, harmonic progressions, dissonant patterns, cadences, etc.), using the 2 theoretical models in the **Music Analysis** module. **Meetups KG** describes 74K+ historical meetups from c.1800 to 1945, mentioning 51K+ people from 5K+ places in Europe. It uses the **Meetups** ontology and is extracted from 1K artists' biographies on open-access digital sources. **Bells KG** describes 88 bells catalogued by the Italian Ministry of Culture¹⁹. It relies on the **Bells** module and is part of the ArCo KG – the largest Italian cultural heritage KG from the Italian General Catalogue of Cultural Heritage. **MusicBO KG** is built via text-to-KG methods [42] on a collection of 137 documents²⁰ on performances and encounters between musicians, composers, and critics happened in Bologna from the 17th century. As mentioned in Section 4, the KG²¹ is used as input to the *bottom-up* modelling of the MusicBO ontology.

5.2 Survey of interest for future applications

To gather interest of adoption, we conducted an online survey in which we ask potential adopters 14 questions regarding their background, relevance, and interest in using music ontologies. The survey was conducted via Google Forms, and distributed in the Semantic Web (SW), Music Information Retrieval (MIR), and Digital Humanities mailing lists – gathering a total of $N = 40$ responses. Among our respondents, 16 work in Semantic Web, 14 in MIR, 16 in Musicology. Most of them have encountered the need for modelling music related data and resources with ontologies (69.3%), focusing primarily on music metadata (25) theory and notation (16), annotations (17) and instruments (17); with 81% doing research or project work related to music data with multiple stakeholders.

Participants were asked to quantify the agreement with statements from 1 (absolutely disagree) to 5 (absolutely agree), 3 being a neutral response (NR; neither agree nor disagree). Results are illustrated in Figure 4. From questions 6-14 we found that: 46.1% find the reuse of existing music ontologies to be challenging (with 41% NR), and the same can be said about the interoperability of existing ontologies (59%; 31% NR), their lack of coverage of concepts related to music history and music cultural heritage (61.5%; 25.6% NR); and the lack of large datasets of competency questions for this domain (63%; 31.6% NR). We also find strong evidence for potential reuse of PON, as participants would be

¹⁶ <https://polifonia.disi.unibo.it/choco/sparql>

¹⁷ <https://polifonia.disi.unibo.it/harmory/sparql>

¹⁸ <https://data-iremum.huma-num.fr/sparql>

¹⁹ <https://dati.cultura.gov.it/sparql>

²⁰ <https://doi.org/10.5281/zenodo.6672165>

²¹ <https://polifonia.disi.unibo.it/musicbo/sparql>

interested in using ontologies for music metadata (79.5%), sources (82%), musical instruments (72%), and music content (59%; 23% NR), as well as a CQ dataset for musical heritage (64%; 33% NR).

5.3 Adoption by Polifonia Stakeholders

In addition to internal and potential adopters, industrial and institutional stakeholders in the Polifonia Stakeholder Network have also expressed interest to use PON resources. These include the **Digital Music Observatory**, concerning the use of **Music Meta** and **Source** to annotate the numerous music resources of the consortium; and the **Université Catholique de Louvain** where Anne-Emmanuelle Ceulemans uses the **Music Analysis** module for studying the annotation of cadences in Josquin des Prez (composer of High Renaissance music).

We have also planned work with **Deezer**, **Songfacts**, and **MusicID** for the evaluation, extension, and reuse of **Music Meta** driven by their resources; and collaborations with the EU H2020 **MuseIT**²² project to extend the ChoCo KG.

6 Availability, sustainability, and FAIRness

PON namespaces are introduced in Section 4, and permanent URIs were created with the W3C Permanent Identifier Community Group. PON is under version control on public GitHub repositories (c.f. Table 1), and all repositories are also published on Zenodo (with associated DOIs) under the CC-BY 4.0 licence. The storage of all resources on GitHub guarantees their persistence beyond the project, with the University of Bologna and the Italian Ministry of Culture (MiC) committed to host and maintain PON on the long term. We also remark that PON is reused as a sibling ontology project of ArCo by MiC [12].

7 Conclusions

This article addresses the creation of the Polifonia Ontology Network (PON), a collection of *expressive* ontologies for musical cultural heritage that enable *interoperability* of existent semantic models for music content and contexts (MusicOWL, Music Notation, Music Ontology, DOREMUS, etc.). Departing from multidisciplinary and domain-expert based requirements in the Polifonia project, we applied and extended the XD methodology for ontology engineering both methodologically (persona/story framework) and technologically (IDEA framework for NLP-assisted co-design). We release 15 new ontologies in PON (v1.0) and the PolifoniaCQ dataset of 361 competency questions under CC-BY 4.0; and provide strong evidence of current and potential reuse by institutional and industrial stakeholders. As next steps, we are planning to perform an extensive CQ-driven evaluation of PON modules; and support our stakeholders and early adopters in the reuse, extension and long term maintenance of our resources.

²² <https://www.muse-it.eu/>



Fig. 4. Selection of questions 4, 6, 7-14 from the online survey, where responses are expressed on a Likert scale ranging from 1 (*strongly disagree*) to 5 (*strongly agree*).

References

1. Achichi, M., Bailly, R., Cecconi, C., Destandau, M., Todorov, K., Troncy, R.: Doremus: Doing reusable musical data. In: ISWC: International Semantic Web Conference (2015)
2. Allik, A., Fazekas, G., Sandler, M.B.: 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)
3. Bechhofer, S., Dixon, S., Fazekas, G., Wilmering, T., Page, K.: Computational analysis of the live music archive. In: ISMIR2014, 15th International Society for Music Information Retrieval Conference (2014)
4. de Berardinis, J., Meroño-Peñuela, A., Poltronieri, A., Presutti, V.: The music annotation pattern. In: Svátek, V., Carriero, V.A., Poveda-Villalón, M., Kindermann, C., Zhou, L. (eds.) Proceedings of the 13th Workshop on Ontology Design and Patterns (WOP 2022) co-located with the 21th International Semantic Web Conference (ISWC 2022), Online, October 24, 2022. CEUR Workshop Proceedings, vol. 3352. CEUR-WS.org (2022), <https://ceur-ws.org/Vol-3352/pattern2.pdf>
5. de Berardinis, J., Meroño-Peñuela, A., Poltronieri, A., Presutti, V.: The harmonic memory: a knowledge graph of harmonic patterns as a trustworthy framework for computational creativity. In: Proceedings of the ACM Web Conference 2023. pp. 3873–3882 (2023)
6. de Berardinis, J., Meroño-Peñuela, A., Poltronieri, A., Presutti, V.: Choco: a chord corpus and a data transformation workflow for musical harmony knowledge graphs. In: Manuscript under review (2023)
7. Blomqvist, E., Hammar, K., Presutti, V.: Engineering ontologies with patterns - the extreme design methodology. In: Ontology Engineering with Ontology Design Patterns - Foundations and Applications, Studies on the Semantic Web, vol. 25. IOS Press, Amsterdam (2016). <https://doi.org/10.3233/978-1-61499-676-7-23>
8. Blomqvist, E., Presutti, V., Daga, E., Gangemi, A.: Experimenting with extreme design. In: Knowledge Engineering and Management by the Masses. EKAW 2010. vol. 6317, pp. 120–134. Springer, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16438-5_9
9. Bottini, T., Carriero, V.A., Carvalho, J., Cathé, P., Ciroku, F., Daga, E., Daquino, M., Davy-Rigaux, A., Guillotel-Nothmann, Gurrieri, M., van Kemenade, P., Marzi, E., Meroño Peñuelala, A., Mulholland, P., Musumeci, E., Presutti, V., Scharnhorst, A.: D1.1 roadmap and pilot requirements 1st version. Tech. rep., EU Commission, The Polifonia consortium (2021)
10. Bruseker, G., Carboni, N., Guillem, A.: Cultural heritage data management: the role of formal ontology and cidoc crm. Heritage and archaeology in the digital age: acquisition, curation, and dissemination of spatial cultural heritage data pp. 93–131 (2017)
11. Carriero, V.A., Ciroku, F., de Berardinis, J., Pandiani, D.S.M., Meroño-Peñuela, A., Poltronieri, A., Presutti, V.: Semantic integration of mir datasets with the polifonia ontology network. In: Proc. of the International Society for Music Information Retrieval Conference, ISMIR (2021)
12. Carriero, V.A., Gangemi, A., Mancinelli, M.L., Marinucci, L., Nuzzolese, A.G., Presutti, V., Veninata, C.: Arco: The italian cultural heritage knowledge graph. In: The Semantic Web–ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part II 18. pp. 36–52. Springer (2019)

13. Carriero, V.A., Gangemi, A., Mancinelli, M.L., Nuzzolese, A.G., Presutti, V., Veninata, C.: Pattern-based design applied to cultural heritage knowledge graphs. *Semantic Web* **12**(2), 313–357 (2021). <https://doi.org/10.3233/SW-200422>, <https://doi.org/10.3233/SW-200422>
14. Carriero, V.A., et al.: The landscape of ontology reuse approaches. *Applications and Practices in Ontology Design, Extraction, and Reasoning* **49**, 21 (2020)
15. Cherfi, S.S., Guillotel, C., Hamdi, F., Rigaux, P., Travers, N.: Ontology-based annotation of music scores. In: Corcho, Ó., Janowicz, K., Rizzo, G., Tiddi, I., Garijo, D. (eds.) *Proceedings of the Knowledge Capture Conference, K-CAP 2017*, Austin, TX, USA, December 4-6, 2017. pp. 10:1–10:4. ACM (2017). <https://doi.org/10.1145/3148011.3148038>, <https://doi.org/10.1145/3148011.3148038>
16. Cherfi, S.S.s., Guillotel, C., Hamdi, F., Rigaux, P., Travers, N.: Ontology-based annotation of music scores. In: *Proceedings of the Knowledge Capture Conference. K-CAP 2017*, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3148011.3148038>, <https://doi.org/10.1145/3148011.3148038>
17. Church, K.W.: Word2vec. *Natural Language Engineering* **23**(1), 155–162 (2017)
18. Diamantini, C., Potena, D., Storti, E.: Ontology-driven kdd process composition. In: *Advances in Intelligent Data Analysis VIII: 8th International Symposium on Intelligent Data Analysis, IDA 2009*, Lyon, France, August 31-September 2, 2009. *Proceedings 8*. pp. 285–296. Springer (2009)
19. Falco, R., Gangemi, A., Peroni, S., Shotton, D., Vitali, F.: Modelling owl ontologies with graffoo. In: *The Semantic Web: ESWC 2014 Satellite Events: ESWC 2014 Satellite Events*, Anissaras, Crete, Greece, May 25-29, 2014, *Revised Selected Papers 11*. pp. 320–325. Springer (2014)
20. Fazekas, G., Sandler, M.B.: 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)
21. Fazekas, G., Raimond, Y., Jacobson, K., Sandler, M.: An overview of semantic web activities in the omras2 project. *Journal of New Music Research* **39** (12 2010). <https://doi.org/10.1080/09298215.2010.536555>
22. Gangemi, A., Peroni, S.: The Information Realization Pattern. In: Hitzler, P., Gangemi, A., Janowicz, K., Krisnadhi, A., Presutti, V. (eds.) *Ontology Engineering with Ontology Design Patterns - Foundations and Applications, Studies on the Semantic Web*, vol. 25, pp. 299–312. IOS Press (2016). <https://doi.org/10.3233/978-1-61499-676-7-299>, <https://doi.org/10.3233/978-1-61499-676-7-299>
23. Gangemi, A., Presutti, V.: Ontology design patterns. In: Staab, S., Studer, R. (eds.) *Handbook on Ontologies*, pp. 221–243. *International Handbooks on Information Systems*, Springer, Berlin, Heidelberg (2009). https://doi.org/10.1007/978-3-540-92673-3_10
24. Gruen, D., Rauch, T., Redpath, S., Ruettinger, S.: The use of stories in user experience design. *International Journal of Human-Computer Interaction* **14**(3-4), 503–534 (2002)
25. Gruninger, M., Fox, M.S.: The role of competency questions in enterprise engineering. In: *Benchmarking — Theory and Practice. IFIP Advances in Information and Communication Technology*. pp. 83–95. Springer, Boston, MA (1994)
26. Guillotel-Nothmann, C.: Knowledge extraction and modelling in the project thesaurus musicarum germanicarum. In: *Prague DH Workshops Session III: Editions* (2020)

27. Guillotel-Nothmann, C., Ceulemans, A.E.: Das diatonisch-chromatische system zur zeit des michael praetorius. eine digitale neuerschließung des syntagma musicum (1619) in verbindung mit dem tanzyklus terpsichore (1612). In: Musik im Umbruch. Michael Praetorius zum 400. Todestag, Harrassowitz (2021)
28. Harley, N., Wiggins, G.: An ontology for abstract, hierarchical music representation. In: Demo at the 16th International Society for Music Information Retrieval Conference (ISMIR 2015), Malaga, Spain (2015)
29. Hart, M.: Preserving our musical heritage: A musician's outreach to audio engineers. *Journal of the Audio Engineering Society* **49**(7/8), 667–670 (2001)
30. Hentschel, J., Moss, F.C., McLeod, A., Neuwirth, M., Rohrmeier, M.: Towards a Unified Model of Chords in Western Harmony. In: Münnich, S., Rizo, D. (eds.) Music Encoding Conference Proceedings 2021. pp. 143–149. Humanities Commons (2022). <https://doi.org/10.17613/4crx-fr36>
31. Hitzler, P., Gangemi, A., Janowicz, K., Krisnadhi, A., Presutti, V. (eds.): Ontology Engineering with Ontology Design Patterns - Foundations and Applications, Studies on the Semantic Web, vol. 25. IOS Press, Amsterdam (2016)
32. Humphrey, E., Salamon, J., Nieto, O., Forsyth, J., Bittner, R., Bello, J.: Jams: A json annotated music specification for reproducible mir research. pp. 591–596 (2014), 15th International Society for Music Information Retrieval Conference, ISMIR 2014 ; Conference date: 27-10-2014 Through 31-10-2014
33. Isaac, A., Haslhofer, B.: Europeana linked open data–data. europeana. eu. *Semantic Web* **4**(3), 291–297 (2013)
34. Jacobson, K., Raimond, Y., Sandler, M.B.: 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)
35. Jones, J., de Siqueira Braga, D., Tertuliano, K., Kauppinen, T.: Musicowl: The music score ontology. In: Proceedings of the International Conference on Web Intelligence. pp. 1222–1229 (2017)
36. KCL, U., NISV, I., KCL, A.M.P., de Berardinis, J., Carriero, V.A., Wigham, M., Poltronieri, A., Ciroku, F., Guillotel-Nothmann, C., Rigaux, P.: D2.1: Ontology-based knowledge graphs for music objects (v1.0) (2021)
37. Kim, H.H.: A semantically enhanced tag-based music recommendation using emotion ontology. In: Intelligent Information and Database Systems. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
38. Kolozali, S., Barthet, M., Fazekas, G., Sandler, M.B.: Knowledge representation issues in musical instrument ontology design. In: ISMIR. pp. 465–470 (2011)
39. Lisena, P., Todorov, K., Cecconi, C., Leresche, F., Canno, I., Puyrenier, F., Voisin, M., Le Meur, T., Troncy, R.: Controlled vocabularies for music metadata. In: ISMIR: International Society for Music Information Retrieval (2018)
40. Lisena, P., Troncy, R.: 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. CEUR Workshop Proceedings, vol. 2065. CEUR-WS.org (2017)
41. Marsden, A.: Music analysis by computer: Ontology and epistemology. In: Meredith, D. (ed.) Computational Music Analysis, pp. 3–28. Springer (2016). https://doi.org/10.1007/978-3-319-25931-4_1, https://doi.org/10.1007/978-3-319-25931-4_1

42. Meloni, A., Reforgiato Recupero, D., Gangemi, A.: AMR2FRED, A Tool for Translating Abstract Meaning Representation to Motif-Based Linguistic Knowledge Graphs. In: Blomqvist, E., Hose, K., Paulheim, H., Ławrynowicz, A., Ciravegna, F., Hartig, O. (eds.) *The Semantic Web: ESWC 2017 Satellite Events*. pp. 43–47. Springer International Publishing, Portorož, Slovenia (2017)
43. Meroño-Peñuela, A., Hoekstra, R., Gangemi, A., Bloem, P., de Valk, R., Stringer, B., Janssen, B., de Boer, V., Allik, A., Schlobach, S., Page, K.: The midi linked data cloud. In: *The Semantic Web – ISWC 2017*. Springer International Publishing, Cham (2017)
44. Moraitou, E., Christodoulou, Y., Caridakis, G.: Semantic models and services for conservation and restoration of cultural heritage: A comprehensive survey. *Semantic Web (Preprint)*, 1–31
45. Poltronieri, A., Gangemi, A.: The music note ontology. In: Hammar, K., Shimizu, C., Küçük McGinty, H., Asprino, L., Carriero, V.A. (eds.) *Proceedings of the 12th Workshop on Ontology Design and Patterns (WOP 2021)*, Online, October 24, 2021. (11 2021)
46. Raimond, Y., Abdallah, S., Sandler, M., Giasson, F.: The music ontology. In: *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007)*. Vienna, Austria (Sep 2007)
47. Rashid, S.M., De Roure, D., McGuinness, D.L.: A music theory ontology. In: *Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music*. p. 6–14. SAAM '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3243907.3243913>
48. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019)
49. Rho, S., Song, S., Hwang, E., Kim, M.: Comus: Ontological and rule-based reasoning for music recommendation system. In: *Advances in Knowledge Discovery and Data Mining*. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
50. Riley, J.: Application of the Functional Requirements for Bibliographic Records (FRBR) to Music. In: *ISMIR*. pp. 439–444 (2008)
51. Smilkov, D., Thorat, N., Nicholson, C., Reif, E., Viégas, F.B., Wattenberg, M.: Embedding projector: Interactive visualization and interpretation of embeddings. *arXiv preprint arXiv:1611.05469* (2016)
52. Szeto, K.: Ontology for voice, instruments, and ensembles (onvie): Revisiting the medium of performance concept for enhanced discoverability. *Code4Lib Journal (54)* (2022)
53. Thalmann, F., Carrillo, A.P., Fazekas, G., Wiggins, G.A., Sandler, M.: The mobile audio ontology: Experiencing dynamic music objects on mobile devices. In: *2016 IEEE Tenth International Conference on Semantic Computing (ICSC)*. pp. 47–54. IEEE (2016)
54. Turchet, L., Antoniazzi, F., Viola, F., Giunchiglia, F., Fazekas, G.: The internet of musical things ontology. *Journal of Web Semantics* **60**, 100548 (2020)
55. Turchet, L., Bouquet, P., Molinari, A., Fazekas, G.: The smart musical instruments ontology. *Journal of Web Semantics* **72**, 100687 (2022). <https://doi.org/https://doi.org/10.1016/j.websem.2021.100687>, <https://www.sciencedirect.com/science/article/pii/S1570826821000573>
56. Viljanen, K., Tuominen, J., Hyvönen, E.: Ontology libraries for production use: The finnish ontology library service onki. In: *The Semantic Web: Research and Applications: 6th European Semantic Web Conference, ESWC 2009 Heraklion, Crete, Greece, May 31–June 4, 2009 Proceedings 6*. pp. 781–795. Springer (2009)

57. Wilmering, T., Fazekas, G., Sandler, M.B.: 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)
58. Zanoni, M., Setragno, F., Sarti, A., et al.: The violin ontology. In: Proc. of the 9th Conference on Interdisciplinary Musicology (CIM14). Citeseer (2014)