

Towards a Knowledge Graph Representation of FAIR Music Content for Exploration and Analysis

Emanuele Storti

Dipartimento di Ingegneria dell'Informazione, via Breccie Bianche 60121, Ancona, Italy

Abstract

This paper introduces the ontological model for a FAIR digital library of music documents which takes into account a variety of music-related information, among which editorial information on documents and their production workflow as well as the score content and licensing information. The model is complemented with annotations (e.g. comments, fingering) on music documents produced by end-users, capable to add a social layer over the framework which enables the building of user-centric music applications. As a result, a machine-understandable knowledge graph of music content is defined, which can be queried, navigated and explored. On top of this, novel applications could be designed, like semantic workplaces where music scholars and musicians can find, analyse, compare, annotate and manipulate musical objects.

Keywords

music score, FAIR Data, Linked Data, Knowledge Graph,

1. Introduction


Digital repositories for musical content have long been used as systems to categorize information on documents related to the musical domain. While some of them only act as metadata catalogs for documents stored in physical libraries, others also host digital versions of the corresponding documents. For instance, the International Music Score Library Project (IMSLP)¹ and the Sheet Music Consortium² describe digitized music documents which typically originate from printed sources, e.g. in the form of scanned images, fully encoded scores or other formats. On the other hand, repositories like MusicBrainz³ are focused on storing information on music production, including metadata on records, artists, performers, and relations among them and to external vendors.

Often through collaborative efforts, some of such repositories have reached significant dimensions and now include up to millions of documents. However, they mostly operate as information silos, storing only a particular kind of information with customised access

TPDL2022: 26th International Conference on Theory and Practice of Digital Libraries, 20-23 September 2022, Padua, Italy

 e.storti@univpm.it (E. Storti)

 0000-0001-5966-6921 (E. Storti)

 © 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

¹<http://imslp.org/>

²<https://digital.library.ucla.edu/sheetmusic/>

³<https://musicbrainz.org/>

rules and data/metadata representation models. Even in case APIs are enabled to access metadata, e.g. for IMSLP and MusicBrainz, they are defined through customized interfaces which do not refer to commonly used standards.

While standard formats for digital representation of scores (e.g., MIDI or MusicXML) have been proposed and are widely adopted by several communities, the current lack of shared solutions for organizing data in musical repositories, like standardized vocabularies, has a great impact on interoperability among different sources and data exchange. As such, in order to integrate disparate sources of information, a manual alignment of heterogeneous datasets needs to be performed on a case-by-case basis, which is an error-prone and a time consuming solution.

Unlike other application domains, the musical domain is witnessing only recently a number of proposals for the definition of uniform vocabularies for metadata, also pushed by international efforts towards Open and FAIR (Findable, Accessible, Interoperable, Reusable) Data [1], especially for libraries maintained by public institutions and foundations. Publishing according to the FAIR principles, which are becoming a requirement for public funding in many countries, means to assign each document a unique digital identifier (DOI), providing a rich set of metadata (Findability), granting the access to (meta)data through common protocols, possibly with authentication and authorisation mechanisms (Accessibility), referring to widely adopted standards, possibly expressed through formal formats which are machine-understandable (Interoperability), and ensuring reuse by declaring the licence and the provenance (Reusability). In this sense, the principles can refer both to metadata (e.g., the description of a score in terms of author, title, date of publication, publisher, number of parts, tonality) and data (the score symbolic content in terms of parts, measures, notes).

As also argued by several authors (e.g. [2, 3]), these principles are meant to support the transition towards a more interconnected and open Web of musical Data, capable to empower both user and machines to more easily retrieve musical resources and, whenever possible, combine them to produce integrated views over disparate datasets, derivative works and innovative applications. Among them, smarter and more flexible search engines capable to query both the metadata of a score, its symbolic content, together with its publishing record, e.g. to retrieve the editorial details of scores composed by “Johann Sebastian Bach” between 1723 and 1750, in the key of “D major”, which include a violin part starting with a given rhythmic pattern, i.e. 16 semi-quavers. On top of this, novel applications could be designed, like semantic workplaces where music scholars and musicians can find, analyse, compare, annotate and manipulate musical objects.

Towards this objective, this paper reports on ongoing work on the development of a FAIR representation of metadata and data related to music digital content. The proposed model, in the form of a knowledge graph, integrates existing standards and ontologies for the representation of metadata on music objects (work, scores, records) and the content of music scores to provide a homogeneous view over disparate data sets. As a result, a machine-understandable graph of music content is defined, which can be queried, navigated and explored.

The rest of this work is structured as follows: Section 2 summarizes relevant work in the Literature focusing on semantic representation of musical content. In Section 3 an

integrated model is proposed, on top of which queries integrating various information can be run, as exemplified in Section 4. Finally, Section 5 concludes the work and discusses future work.

2. Related work

The use of different representation mechanisms, file formats and schemas for the large mass of available documents on the web, bring interoperability issues that make integration of data challenging. To overcome such shortcomings, semantic technologies have been exploited to define vocabularies, taxonomies or ontologies providing the terminology that can be used to annotate documents. Represented in a formal and unambiguous format, such models enable the definition of machine-readable descriptions and ultimately the representation of knowledge in a processable way. In this context, the term “Linked Data” refers to a set of semantic technologies and publication practices that are used to create a graph of interconnected datasets. A distinguished principle of this approach is that each data element has a unique identifier (URI) over the web that can be reused by other datasets. To make an example, in DBPedia⁴, a project aiming to extract structured content from Wikipedia and publish them as Linked Data, the URI “https://dbpedia.org/resource/Johann_Sebastian_Bach” represents Johann Sebastian Bach. If digital libraries reuse such an URI to refer to Bach, instead of redefining a custom identifier, their integration would be much facilitated. Furthermore, according to the Linked Data principles, datasets must be accessible through standard protocols such as HTTP and must be represented through standard and self-documenting languages like RDF [4]. This language enables to represent information as a set of basic statements (or triples) in the form subject-property-object. The union of the triples generates a so-called knowledge graph. Finally, data can be queried at Web scale through the SPARQL language. In the following, some ontologies for the representation of music metadata and content are summarized.

Music ontology [5] is a modular and extensible ontology to formally represent music-related information, that has been adopted by several projects including BBC Music and DBTune. Its main purpose is to provide the terminology to interlink different online catalogues. While a basic level of detail only deals with purely editorial information, a second level introduces the concept of event, which is used to describe a workflow involving the composition of a musical work, its arrangement, performances of such an arrangement and recordings of the performances. Music ontology builds on FOAF, a vocabulary for describing people, groups of people and organisations, on the event ontology, a vocabulary for describing events, on the timeline ontology and the Functional Requirements for Bibliographic Record (FRBR) (discussed in Section 3).

Several formats have been proposed for the symbolic content of music score. Among them, MIDI⁵, originally presented in 1981, has been the most popular technical standards for a communication protocol and digital interface enabling a variety of digital system

⁴<http://dbpedia.org>

⁵<https://www.midi.org/specifications/file-format-specifications>



Figure 1: Incipit of the Brandenburg Concerto No.5 in D major, Johann Sebastian Bach.

to record, edit and play music. While MIDI is more focused on connectivity and music playback than to actually representing symbolic content, more recently MusicXML[6] has been proposed as an XML-based format for encoding western musical notation. It is intended for the exchange of music documents across different scorewriters and other applications. The Music Encoding Initiative (MEI)⁶ is a community-driven, open-source effort to define a system for encoding musical documents in a machine-readable structure. Like MusicXML, MEI is encoded as an XML language, but includes a more advanced representation of notations beyond the common western one (e.g. mensural and medieval neume notations).

More recently, some ontologies have been proposed providing RDF vocabularies to describe the symbolic content of music scores. The MIDI Linked Data Cloud [7] proposes to use the Linked Data approach to interconnect symbolic music descriptions contained in MIDI files, while MusicOWL [8] is an ontology including classes and properties to fully represent a MusicXML score in RDF.

Similarly, in [9] a framework is proposed for extracting knowledge from music scores that can be inferred from music notation, e.g. phrases, cadences, dissonances. With a different objective, the Music Theory Ontology [10] aims at defining basic theoretical musical concepts to build a model useful for music education and analysis.

On top of the mentioned approaches, some frameworks have been proposed to support specific applications for analysis. As an extension of the MIDI Linked Data Cloud, the HaMSE ontology [11] is devised to support musicological analysis by harmonizing different representations (audio and score with a mutual alignment), and by including musicological features such as chord progressions, rhythmic patterns or intervals. With a focus on annotation of music performances, in [12] the MELD framework is introduced, augmenting the MEI-encoded score elements with real-time annotation of a score during a performance. Following an approach that is complementary to the present work, the Audio Commons Ontology [13] builds on the Music Ontology for the representation of audio content in the broader context of audio production and sharing, following an approach towards the interoperability of different repositories.

⁶<https://music-encoding.org/>

3. Semantic model of the framework

This section is devoted to introduce the main components of the model used for the representation of information, in the context of an online digital library of music documents. The model takes into account editorial information on documents and their production workflow as well as the score content and licensing information. The model is complemented with annotations (e.g. comments, fingering) on music documents produced by end-users, capable to add a social layer over the framework which enables the building of user-centric music applications. In order to be fully compliant with the FAIR principles, both metadata and data are represented by referring to open shared vocabularies represented in the RDF language. Hence, the final model stems from the integration of such standards and ontologies, as also reported in the following subsections:

- generic metadata on documents and other resources are expressed through Dublin Core properties;
- metadata on music works, scores, recordings, performances and the workflow for their creation are represented through the Music Ontology;
- information on provenance is represented through the PROV-O ontology;
- the content of music scores is represented through the MusicOWL ontology;
- licensing information is represented through the Music Ontology and the Creative Commons schema;
- information related to user content, which can be attached to any music information in the model is represented through the Web Annotation Vocabulary.

To avoid redefining specific terms to express values of metadata, URIs from a number of further external resources have been reused. They include DBPedia and Wikidata⁷ for artists' names, the Tonality ontology⁸ to represent the tonalities (e.g. E minor), the Music Vocabulary⁹ which includes a taxonomy of music forms (e.g. Concerto, Sonata) and a taxonomy of ensemble types (e.g., Ensemble, Orchestra).

3.1. Core document metadata

Dublin Core (DC)[14] is a set of 15 metadata properties for describing generic resources on the web, either physical or digital, formulated by the Dublin Core Metadata Initiative (DCMI). In the context of Linked Data, Dublin Core is one of the most popular vocabularies in RDF and is extensively used for resource description and to provide interoperability for metadata vocabularies among different datasets in a variety of domains.

Among the metadata properties, a resource can be characterized in terms of a title, a creator, one or more subjects useful for descriptive or classification purposes, a textual description, a publisher, a publication date, a type, a format, a source to specify one or more resources from which the resource is derived, the language, specification of rights held in and over the resource. Other properties are defined by DCMI Metadata Terms,

⁷<https://www.wikidata.org/wiki/>

⁸<http://purl.org/ontology/tonality/key/>

⁹<http://www.kanzaki.com/ns/music>

which extends Dublin Core with further terms, e.g., to specify that a resource isPartOf another resource, versioning information, the licence and the rightsHolder among others. An example of music digital library using some Dublin Core properties is the Sheet Music Consortium.

3.2. Document production workflow

Aspects related to music production are defined through classes and relations from the Music Ontology. Like several other ontologies focusing on the representation of musical catalogs, Music Ontology is built on top of a generic and flexible model named Functional Requirements for Bibliographic Record (FRBR), proposed by the International Federation of Library Association (IFLA). The model is aimed to describe documents and their evolution and is particularly suited for both physical and digital resources. The representation of a musical object is done in FRBR at various levels of abstractions, from the generic concept to the specific realization, through the following main elements:

- Work is an abstract concept representing an artistic creation, independently of its concrete realizations, e.g. the Brandenburg Concerto No.5 by J. S. Bach.
- An Expression is the realisation of the artistic content of a Work. For instance, each version of the score for Bach's concerto that has been published is a different expression of the same work. It can be realized through one or more Manifestations.
- A Manifestation represents a particular physical or electronic embodiment of an expression, e.g. the specific formats in which a particular edition of a score can be available: in textual form, as a scanned PDF, in MusicXML. A Manifestation is exemplified by one or more Items.
- An Item is a particular instance of a Manifestation, for instance a specific copy of a record.

The Music Ontology interconnects such elements through the notion of event, e.g. a Composition is an event made by a MusicArtist producing a MusicalWork. This last represents an abstract entity and not a particular concrete realization of it (e.g., a published score or a recording). An Arrangement is an event which produces a score out of a work. A Performance produces a Sound which can be recorded. A Recording event takes a work as input and produces a Signal which can then be published as a Record. A PublishedScore represents a concrete score (i.e., a manifestation), which has a title, a licence, a publication date, a publisher, and may be available in different formats.

Dublin Core properties have been used to specify the title of a musical work or a score, the composition or recording date, the format of a score (as a MIME type), while possible derivation from other resources are represented through the PROV-O¹⁰ property wasDerivedFrom.

The following code shows a fragment of the RDF triples representing the metadata for the published score of the "Brandenburg Concerto No.5" by Johann Sebastian Bach, with a title and a date related to its creation, a public domain licence, two publishers

¹⁰<https://www.w3.org/TR/prov-o/>

and Leipzig as a publishing location. The namespace before the URIs are shorthand for the full URI namespace of the corresponding ontology (“mo” stands for the full Music Ontology namespace, “cc” for the Creative Commons schema, “dc” for Dublin Core, “dbpedia” for DBpedia). Please note that, whenever possible, the value of properties are URIs taken from external sources. In some cases, e.g. the names of the publishers, or the title, a simple string (Literal) was used.

```
:pscoreBrConcert5 rdf:type mo:PublishedScore;
  mo:licence cc:PublicDomain.
  dc:title "Brandenburg Concerto No.5 in D Major";
  dc:date "1851"^^xsd:gYear;
  mo:publisher [ a foaf:Agent;
    rdfs:isDefinedBy dbpedia:Bach_Gesellschaft];
  mo:publisher [ a foaf:Agent;
    rdfs:isDefinedBy dbpedia:Breitkopf_&_Härtel];
  mo:publishing_location dbpedia:Leipzig.
```

3.3. Score content

The content of a score is represented through the MusicOWL ontology, which provides the terminology for the RDF representation of MusicXML documents. This format was chosen over others because MusicXML is one of the most popular and widely adopted file format for music content encoding and sharing. Indeed, several scorewriter applications include import/export tools to/from this format (e.g. Cubase, MuseScore, Finale), making the production and the sharing of MusicXML files easier, and several digital libraries provide MusicXML file sources¹¹.

The ontology includes classes to represent one or more ScorePart, which includes a Staff, which in turn has Voices. A score part include a set of Measures, which can have multiple NoteSets, i.e. containers of notes. A Note is characterized by a Duration (e.g. 1/4) and specifies a natural value (the pitch) and a possible modifier (e.g., sharp, flat, double sharp). A fragment of the beginning of the lead violin part for the Brandenburg Concerto No.5 (see also Figure 1) is represented as follows:

```
:measure1 a mso:Measure;
  mso:hasNoteSet :noteset1.

:noteset1 a mso:NoteSet;
  mso:hasNote :note1, :note2, :note3, :note4;
  mso:duration mso:Quarter;
  mso:nextNoteSet :noteset2.

:note1 a chord>Note;
  chord:natural note:D.

:note2 a chord>Note;
  chord:natural note:D.
```

¹¹A partial list is available at <https://www.musicxml.com/music-in-musicxml/>.


```
:note3 a chord:Note;
  chord:natural note:F;
  chord:modifier chord:sharp.
```

```
:note4 a chord:Note;
  chord:natural note:F;
  chord:modifier chord:sharp.
```

3.4. Licensing information

A relevant information for any resource published in a FAIR repository is the specification of the licence, which determines what operations are permitted or prohibited on the resource and what requirements are set for users. Particularly, the FAIR principle R1.1 requires that (meta)data are released with a clear and accessible data usage license. Licensing information for a published score or a record have been described through the `mo:licence` property, specifying the licence by relying on the Creative Commons schema. In such a way, it is possible to declare specific CC licenses by combining different terms, thus making distinct aspects of the licence machine-understandable. The following snippet of RDF assigns a Creative Commons licence named Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) to a published score.

```
<#publishedscore> mo:licence :cc-by-nc-nd-4.
:cc-by-nc-nd-4 a cc:Licence;
  cc:permits cc:Reproduction;
  cc:permits cc:Distribution;
  cc:prohibits cc:CommercialUse;
  cc:requires cc:Attribution;
  rdfs:seeAlso <https://creativecommons.org/licenses/by-nc-nd/4.0/>;
```

3.5. Annotation of scores

Fingering is the process of mapping each note on a score to a fingered position on some instrument. Apart from didactic music notes, fingering indications are typically not reported in published scores, leaving them to execution choices. Nonetheless, especially for music students, finding the most effective fingering may be tricky and often crucial in the educational process. As well as fingering, other comments or notes are very often reported in music scores to annotate information which is relevant for the performer during the educational process or the performance. However, no standard representation of this information has been developed on the topic, also because of its rather informal structure. For this reason, we rely on the Web Annotation Data Model [15] which describes a structured model and format to enable annotations on generic target resources to be shared and reused across different hardware and software platforms. In particular, we refer to classes and properties defined by the Web Annotation Vocabulary¹². The class

¹²<https://www.w3.org/TR/annotation-vocab/#bib-annotation-model>



Figure 2: Fragment of the knowledge graph describing the concept of composition, musical work and score, with a focus on a published score.

oa:Annotation is used to declare an annotation and its metadata, including the creator and the datetime. The annotation is linked to a target and a body. The former is any given musical element in a score, being it a note, a noteset, a measure, a part or the whole score, or others. On the other hand, the latter is the actual content of the annotation, specifying its value, its format as well possible further metadata, e.g. language. In the following, an example is shown of a fingering indication which is created by a user and attached to a note:

```
:note1 a chord:Note ;
      mso:hasOctave "4"^^xsd:int ;
      chord:natural note:D.

:anno1 a oa:Annotation ;
      dcterms:creator :user1;
      dcterms:created "2022-06-15T17:31:00.000"^^xsd:dateTime;
      oa:hasTarget :note1 ;
      oa:hasBody [
        a oa:TextualBody;
        rdf:value "2" ;
        dc:format "text/plain" ] .
```

4. Querying the music content

The representation of the model through a RDF knowledge graph enables the possibility to make queries to extract relevant information, that can involve different aspects of the musical content. For instance, queries on bibliographic information of the document, e.g. about the creator or the publisher, can be combined with information on the production process, performances, as well as with information on the structure of the score, e.g. its parts, its specific melodic/harmonic/rhythmic content. Furthermore, information on user comments and annotations can be integrated as well. Queries are expressed through the SPARQL language, which enables easy data access and interoperability with external applications.

As an example of queries that can be expressed on the graph, the following one asks for published scores authored by “Ludwig van Beethoven” which are released under the public domain licence, for which a PDF version is available.

```
SELECT ?pscore
WHERE {
  ?c mo:composer dbpedia:Ludwig_van_Beethoven.
  ?c mo:produced_work ?work.
  ?work mo:arranged_in ?arr.
  ?arr mo:produced_score ?pscore.
  ?pscore mo:licence cc:PublicDomain.
  ?pscore mo:available_as ?pscore_pdf.
  ?pscore_pdf dc:format ?format.
  FILTER (?format = mime:application/pdf).
}
```

To make a further example, the following query searches for scores including a measure with a quarter and a semiquaver, and extracts comments on the measure.

```
SELECT ?pscore ?measure ?value
WHERE {
  ?pscore mso:movement ?mov.
  ?mov mso:hasScorePart ?part.
  ?part mso:hasMeasure ?measure.
  ?measure mso:hasNoteSet ?ns1.
  ?measure mso:hasNoteSet ?ns2.
  ?ns1 mso:nextNoteSet ?ns2.
  ?ns1 mso:hasDuration ?dur1.
  ?dur1 a mso:Quarter.
  ?ns2 mso:hasDuration ?dur2.
  ?dur2 a mso:16th.
  ?ann oa:hasTarget ?measure.
  ?ann oa:hasBody ?body.
  ?body rdf:value ?value.
}
```

5. Discussion

This paper introduced the ontological model for a FAIR digital library of music documents which takes into account a variety of music-related information. The resulting RDF model has the shape of a knowledge graph, where all the information can be explored and queried according to the Linked Data approach, relying on standard tools and protocols. As an example, Figure 2 shows a graphical representation of a fragment of the knowledge graph describing the “Brandenburg Concerto No. 5” by Johann Sebastian Bach.

This ongoing work represents the roots on which a user-centric framework for documentation, editing and exchange of documents will be built. Future steps will be devoted to develop the application layer on top of the model, which will include graphical user interfaces to enable user-friendly browsing and exploration of the graph, annotation of scores, their analysis and sharing.

On the other hand, several challenges still need to be address that call for extensions of existing model schemas. On the one hand, almost all ontologies for the symbolic representation of scores currently focus only on the western music notation. As a consequence, other music traditions cannot be fully represented by using such formats, as some work point out (e.g. [16]). Furthermore, existing models do not take into account novel notation symbols, which may be characteristic of specific musical instruments and specific music communities, e.g. contemporary compositions often include non-standard notations that can hardly be automatically understood by optical music recognition systems and hence represented through existing models.

References

- [1] 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* 3 (2016) 1–9.
- [2] A. Hofmann, T. Miksa, P. Knees, A. Bakos, H. Sağlam, A. Ahmedaja, B. Yimwadsana, C. Chan, A. Rauber, Enabling FAIR use of Ethnomusicology Data—Through Distributed Repositories, *Linked Data and Music Information Retrieval, Empirical Musicology Review* 16 (2021) 47–64.
- [3] D. M. Weigl, T. Crawford, A. Gkiokas, W. Goebel, E. Gómez, N. F. Gutiérrez, C. C. Liem, P. Santos, FAIR Interconnection and Enrichment of Public-Domain Music Resources on the Web, *Empirical Musicology Review* 16 (2021) 16–33.
- [4] O. Lassila, R. R. Swick, et al., Resource Description Framework (RDF) model and syntax specification (1998).
- [5] Y. Raimond, S. A. Abdallah, M. B. Sandler, F. Giasson, The Music Ontology., in: *ISMIR*, volume 2007, 2007, p. 8th.
- [6] M. Good, MusicXML for notation and analysis, *The virtual score: representation, retrieval, restoration* 12 (2001) 160.
- [7] A. Meroño-Peñuela, R. Hoekstra, A. Gangemi, P. Bloem, R. d. Valk, B. Stringer,

- B. Janssen, V. d. Boer, A. Allik, S. Schlobach, et al., The MIDI Linked Data Cloud, in: International Semantic Web Conference, Springer, 2017, pp. 156–164.
- [8] J. Jones, D. de Siqueira Braga, K. Tertuliano, T. Kauppinen, MusicOWL: the Music Score Ontology, in: Proceedings of the International Conference on Web Intelligence, 2017, pp. 1222–1229.
 - [9] S. S. Cherfi, C. Guillotel, F. Hamdi, P. Rigaux, N. Travers, Ontology-based annotation of music scores, in: Proceedings of the Knowledge Capture Conference, 2017, pp. 1–4.
 - [10] S. M. Rashid, D. De Roure, D. L. McGuinness, A Music Theory Ontology, in: Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music, SAAM '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 6–14.
 - [11] A. Poltronieri, A. Gangemi, The HaMSE Ontology: Using Semantic Technologies to support Music Representation Interoperability and Musicological Analysis, arXiv preprint arXiv:2202.05817 (2022).
 - [12] D. Weigl, K. Page, A framework for distributed semantic annotation of musical score: “Take it to the bridge!”, in: 18th International Society for Music Information Retrieval Conference, International Society for Music Information Retrieval, Suzhou, China, 2017.
 - [13] M. Ceriani, G. Fazekas, Audio Commons Ontology: a data model for an audio content ecosystem, in: International Semantic Web Conference, Springer, 2018, pp. 20–35.
 - [14] S. Weibel, J. Kunze, C. Lagoze, M. Wolf, Dublin Core Metadata for Resource Discovery, Internet Engineering Task Force RFC 2413 (1998) 132.
 - [15] R. Sanderson, P. Ciccicarese, H. Van de Sompel, S. Bradshaw, D. Brickley, L. J. G. Castro, T. Clark, T. Cole, P. Desenne, A. Gerber, et al., Open Annotation Data Model, W3C community draft 8 (2013).
 - [16] P. Proutskova, A. Volk, P. Heidarian, G. Fazekas, et al., From Music Ontology Towards Ethno-Music-Orthology, in: Proceedings of 21st International Society for Music Information Retrieval Conference (ISMIR 2020), ISMIR press, 2020, pp. 923–931.