Representing Complex Knowledge For Exploration and Recommendation: The Case of Classical Music Information

Pasquale Lisena[1][0000−0003−3094−5585] and Raphäel Troncy[1][0000−0003−0457−1436]

EURECOM, Sophia Antipolis, France
{pasquale.lisena, raphael.troncy}@eurecom.fr

Abstract. In Digital Humanities, one of the main challenge consists in capturing the structure of complex information in data models and ontologies, in particular when connections between terms are not trivial. This is typically the case for librarian music data. In this chapter, we provide some good practices for representing complex knowledge using the DOREMUS ontology as an exemplary case. We also show various applications of a Knowledge Graph leveraging on the ontology, ranging from an exploratory search engine, a recommender system and a conversational agent enabling to answer classical music questions.

1 Introduction

The Semantic Web offers languages and technologies for representing information regardless of the fields of human knowledge. Digital Humanities and Digital Libraries in particular found, in the graph structure promoted by the Semantic Web, a suitable framework for describing data about literature, arts, and history, in which the different elements are often linked to each other by a complex set of relationships.

An exemplary case is the one of classical music, which can be described through very rich metadata. Taking as example a well-known masterpiece such as Beethoven's *Moonlight Sonata*, it is possible to describe the music work as composed by the German composer, or its score in the handmade original version or in the different printed editions, the multiple interpretations by pianists and – in case of arrangement – by other instruments. Related to these interpretations, the performances, recordings, music albums can also be described. This media production chain involves several actors: composers, performers with different but well-defined roles, conductors, technicians, etc. Research about classical music requires specific attention strategies [17, 35] which differs from the ones of *popular music*[1], in which pieces are commonly described using only the track name and the artist name. In other words, the representation challenge, which is

[1] With *popular music*, we refer to all those genres that do not fall under the definitions of classical music, jazz or world music, e.g. pop, rock, hip-hop, funk, rap, electronica, dance.

common to any kind of music description, is particularly evident when targeting classical music. *Which model can be successfully applied for better representing complex knowledge such as music information? To which questions should this model be able to give an answer, in order to give benefit to final users and music scholars? How is final user consumption impacted by music specialised Knowledge Base?*

In this chapter, we describe a complete strategy for representing librarian data, with the goal of modelling a specialised Knowledge Graph. We take as example the work realised in the context of the DOREMUS project[2], which published music metadata – music events, resources, and agents – coming from different libraries and archives as Linked Data. However, the same foundations, patterns and directions can be applied to other creative industries.

2 Related work

This section aims to give to the reader an overview of related works in the literature covering the representation of cultural object in general, and music in particular, and the realisation of Digital Libraries as Knowledge Graphs.

2.1 Ontologies for libraries and music data

Different models and ontologies have been proposed for representing music information using Semantic Web Technologies. An important role as foundation for many music and – in general – cultural ontologies is hold by the **Functional Requirements for Bibliographic Records (FRBR)** conceptual model. Published by the International Federation of Library Associations and Institutions (IFLA) for the first time in 1998, this schema defines four distinct states in which a generic cultural object can exists: the *Work* – intended as the artistic or intellectual idea and aim – is realised through a specific set of choices in the content to which we refer as *Expression*; this one comes in the reality in a physical shape, the *Manifestation*, which can be produced in one or more single *Items* (Figure 1). For example, Victor Hugo's story of a hunchback bell-ringer of Notre-Dame Cathedral (*Work*) is formalised in a specific choice in the words which compose *Notre-Dame de Paris* book (*Expression*), which has been published in different editions (*Manifestation*) with a certain number of copies (*Items*). In the same way, Beethoven's idea of a pastoral symphony (*Work*) is realised through the composition of the *Symphony No. 6 in F major, Op. 68* (*Expression*). This expression can be then recorded on a specific edition on CD (*Manifestation*) in several copies (*Items*).

Among the music models relying on FRBR, the **Music Ontology (MO)** [30] is the most known one in the Semantic Web community. This ontology extends the Timeline Ontology and the Event Ontology [28, 29], providing a set of music-specific classes and properties for describing musical works, performances and

[2] http://www.doremus.org/

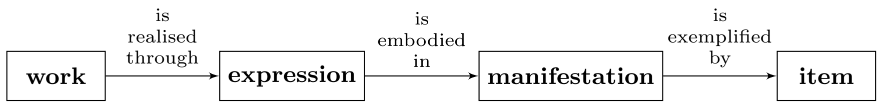

Fig. 1. FRBR diagram

tracks, together with fragments of them. The authors foresee the use of taxonomies and vocabularies for populating the values of certain properties, like keys, instruments and genres. Several examples of interconnecting MO to other datasets, whether they describe music or other kind of data (i.e. DBpedia) are shown in [31]. Even with some attention to classical music – visible in classes and properties like composition, catalogue number, arrangement – MO reveals a strong connection with the track-based vision of the music. Some notable absences that can confirm this statement are alternative roles in composition rather than the sole composer, alternative titles with specific properties (original title, given title, translation), details in the number of foreseen instruments, connections between performers and instruments in a performance, etc. Beside the simplicity of adopting the model, MO is quite far from being able to represent the information coming from specialised classical music archives.

In the domain of cultural objects – to which music belongs – one of the most popular ontologies is **FRBRoo** [10] (where "oo" stands for *object-oriented*), born as a dialog of the librarian FRBR model and the CIDOC-CRM [9]. CIDOC-CRM is an ontology developed for the museum domain. One of its main characteristics is the importance given to events: no object can exists without a specific creation event, and events are required for specifying the object location in a museum or describing its appearance through observation. The harmonisation of FRBR and CIDOC-CRM gives birth to the Work-Expression-Event triplet[3] pattern of FRBRoo (Figure 2): the abstract intention of the author (Work) exists only through an Event (i.e. the composition) that realises it in a distinct series of choices called Expression(s). Considering the book *Moby Dick* as an example, the artistic object takes birth when the idea (Work) of the author Melville are written (Event) in the succession of words (Expression). The relations between these classes and the related subclasses represent one of the strengths of the model thanks to the wide expressiveness gained from this. In FRBRoo, one can link a work with another one (a specific critic edition or the French translation), add more details about the creation event (where and when it took place), add derivatives works (the 1956's movie *Moby Dick*) or works that are components of a complex one (the critics essays contained in a particular edition). Similarly for music, the composition effort of the musician (Event) concretises his/her music idea (Work) into a specific set of choices (Expression) about notes, pauses, dynamics, etc. The work can also here be derived (orchestration, arrangement, variation) or included in other works (movies, theatre).

[3] Not to be confused with an RDF triple.

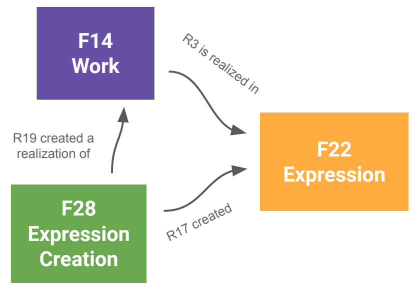

Fig. 2. The triplet pattern in FRBRoo

Other works focus instead on the music content itself, rather than on metadata. An attempt to represent the whole music theory fundamentals lead to the development of the **Music Theory Ontology** [32], with the goal of computing analysis and inference by relying on the music rules. The information contained in a music score is instead represented in the **MusicNote Ontology** [6].

2.2 Digital Libraries as Knowledge Graphs

Semantic Web technologies have a strong predisposition for representing the human knowledge, making it open and accessible for public consumption, and enabling connections between datasets. This predisposition has fed in last decade a new attitude for sharing the knowledge beyond the institutional and national borders, embodied by international consortia like the International Association of Music Libraries, Archives and Documentation Centres (IAML)[4] or in projects like Europeana [14], OpenGlam [11], and datos.bne.es [37] to name a few. The benefits that Semantic Web can offer to Digital Libraries (DL) have been reported by several works, among which the most influential is the study made by the W3C Library Linked Data Incubator Group in [36], that can be summarised as follows:

- it provides methods and standards for integrating different metadata sources, like bibliographic, controlled vocabulary, annotations and non-library sources such as Wikipedia, GeoNames, MusicBrainz, and others;
- it offers solution for interoperability among cultural institutions, promoting the re-use of resources through shared identifiers (uris) and fostering interdisciplinary research;
- it triggers the passage from specific data structures to models whose durability and robustness is ensured by the semantic description of classes and relations;
- it increases the visibility of cultural data on the Web;
- it encourages a discovery approach of cultural information based on navigation on links ("following one's nose");

[4] https://www.iaml.info/

– it opens the librarian knowledge to developers, researches and other communities going beyond library-specific formats;
– it enables advanced use of librarian knowledge, including smart search, reasoning and AI applications.

Accordingly, Semantic Web technologies have gained a central role also on the music domain, that has reached the LOD world. In [4], a traditional music Digital Library is developed through the conversion of metadata in RDF and its enrichment through linking to external Linked Data resources, although the elements in the resulting graph continue to be conceived as separate records instead of interconnected nodes.

Different experiences about converting data from the librarian format MARC to RDF have been explored[5]. The datos.bne.es project developed MARiMbA [37], a software for the conversion of MARC data from the Spanish National Library in RDF, using the FRBR model.

The need for harmonisation of musical metadata coming from different sources and formats led to different technical solutions, often making use of Semantic Web technologies. The **Transforming Musicology** project created **In-Concert** [26], a RDF dataset of performance metadata collected from concert ephemera, such as programmes, reviews, adverts, etc. The dataset has been created by converting and connecting data sources in other formats. A similar workflow made possible the creation of the **JazzCats** dataset[6], specialised in jazz performances [25].

Other semantic music libraries that is worth to mention are the **MIDI Linked Data Cloud** [24, 23], a big archive of MIDI information represented in RDF, and the **Listening Experience Database (LED)** [2], a knowledge base of annotations about music listening. A more complete list about music datasets on the web has been collected in **Musical Data on the Web (musoW)** [8].

3 Modeling a classical music ontology

In this section, we describe our approach that lead to the modeling of the DOREMUS ontology.

3.1 Start and end with competency questions

The design of a new data model (or ontology) may derive from different methodologies, which include, for example, ontology design patterns or mapping of already existing data structures. In some context, the best solution relies on a set of *competency questions*, which the model is expecting to answer [33]. In the case of music metadata, the competency questions help in understanding if existing

[5] https://github.com/search?q=marc2rdf
[6] http://jazzcats.cdhr.anu.edu.au/

ontologies (see Section 2.1) can be re-used or if it is required to realise a new one.

With this background, a list of competency questions have been collected from experts coming from the world of music libraries, radios and concert halls[7]. These questions reflect the real needs of these institutions and reveal problems that they face daily in the task of selecting information from existing databases (e.g. concert organisation or broadcast programming) or for supporting librarian and musicologist studies. They can be related to practical use cases (the search of all the scores that suit a particular formation), to musicologist topics (the music of a certain region in a particular historical period), to interesting stats (the works usually performed or published together), or to curious connections between works, performances or artists. Most of the questions are very specific and complex, so that it is very hard to find their answer by simply querying the search engines currently available on the web.

Some examples of those questions are:

- Give me the list of works composed by Mozart in the last 5 years of his life;
- Give me the works of chamber music that involves at most violin, clarinet and piano, except from the sonatas for violin and piano and clarinet and piano;
- Give me all the works interpreted on at least one instrument different from the casting of the work;
- Give me all the performances in which a composer interprets his or her works;
- Give me the name of the vocal soloist most recorded by Radio France in 2014.

Among them, we can also find questions that overflow the domain, because they contain aspects that go beyond the music information and involve other kind of knowledge. An example is *Retrieve a list of works of chamber music composed in the 19th century by Scandinavian composers*: it requires knowledge of the birth place of the composer, and if this place is located in one of the Scandinavian countries. We can state that these are very interesting questions, because they are the ones that can fully exploit the advantages of Linked Data technologies. In fact, this kind of queries are quite far from having an answer in a traditional relational database systems. The Web of Data gives the possibility of performing federated queries involving the LOD cloud, in particular datasets such as GeoNames [38] or DBpedia [3] (Figure 3). For this reason, the interconnection of the data is crucial.

The inability of other models to deal with some of the information required by competency questions has been crucial for the decision of designing a new ontology, the DOREMUS model, which can better answer to these questions. The competency question will then be used for validating the model and the final knowledge base.

[7] The full list of competency questions is available at https://github.com/DOREMUS-ANR/knowledge-base/tree/master/query-examples

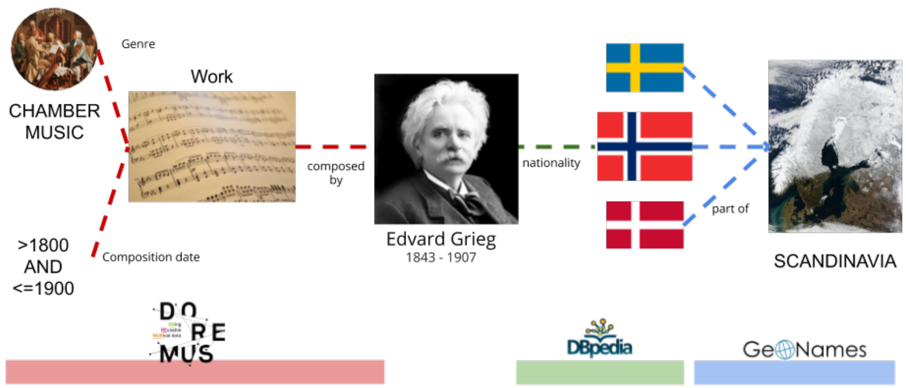

Fig. 3. *Retrieve a list of works of chamber music composed in the 19th century by Scandinavian composers* would require 3 different knowledge bases in order to be answered.

3.2 Designing a music model

A common practice in ontology development [12] consists in relying on already established ontologies and extend them in order to introduce specialised classes and properties. Following this practice, the DOREMUS model is built on top of FRBRoo (see Section 2.1). This choice relies on different motivations:

- **It is a librarian model.** Being popular in librarian archives, FRBRoo appears familiar to cataloguers and fits well with other bibliographic data of music records.
- **It is a bridge to other cultural objects.** The model is ready to be used for describing the interconnection of different arts. FRBRoo provides properties for linking a work such as a musical piece with the poetry that has been adapted in the lyrics or with the film having it in its soundtrack.
- **All triplets are optional.** The Work-Expression-Event pattern ensures that each step of the life of a musical work can be modelled separately, following the triplet structure. For a classic work, we will have a triplet for the composition, one for each performance event, one for every manifestation (e.g., the score), all connected in the graph. Each triplet contains information that at the same time can live autonomously and be linked to the other entities. This provides the freedom of representing, for example, a jazz improvisation as extemporaneous performance not connected to a particular pre-existing work, or to collect all the recordings of a piece of world music.
- **The event expressivity.** In FRBRoo, the creation of a work (physical or performative) can be modelled as a unique event, which in turn is composed of a series of different activities, each one carried by a specific person. In our case, this way of representing the creative process matches perfectly music performances – in which every musician gives a distinct contribution to the sound – or music composition – in which, for example, we can separate the work of the composer and the lyricist.

A modelling group, composed of experts in music cataloguing and knowledge representation, extended FRBRoo, creating new specific classes and properties capable of describing aspects of a work that are related to music, such as the musical key, the genre, the tempo, the instrument, etc. The result of this work is the **DOREMUS Ontology** [7]. Each part of the music production process is considered as an Event that gives birth to a new Work and a new Expression: this leads to the creation of classes like Performance Work or Recording Expression. The graph depicted in Figure 4 shows a real example from our data: Beethoven's *Sonata for piano and cello n.1*.[8]

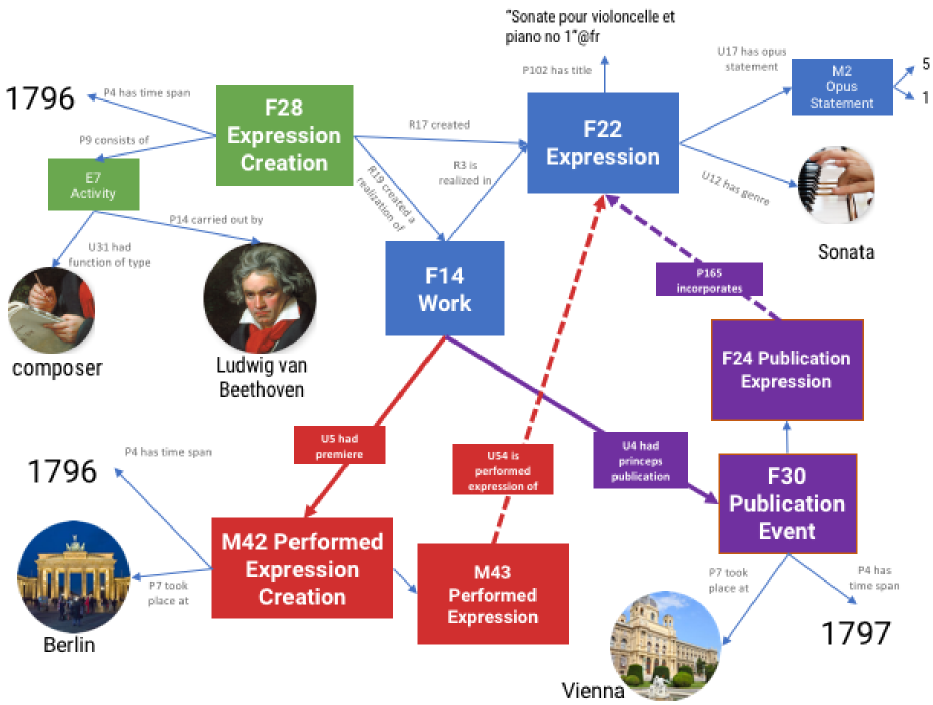

Fig. 4. Schema representing the modelling of Beethoven's *Sonata for cello and piano n.1.* The work and expression (in blue) are linked to the entities representing the composition event (in green) the premiere (in red) and the first publication (in purple).

3.3 Controlled Vocabularies

Describing music is an activity that involves an important number of terms coming from domain-specific glossaries. In addition to the cross-domain concept of genre, we can mention musical keys, instruments or catalogues of compositions.

[8] http://data.doremus.org/expression/614925f2-1da7-39c1-8fb7-4866b1d39fc7

Libraries and musical institutions have different practices for describing this kind of information. In the best case, they make use of thesauri that are often available in different incompatible formats, and that can be either internally defined or standardised by larger communities such as the International Association of Music Libraries (IAML). In other cases, this information is codified in free text fields, delegating to the editors the responsibility of following the living practice about syntax and lexical form.

For the description of music-specific concepts, we published controlled vocabularies (using SKOS and MODS standards), realised and enriched by an editorial process that involved also librarians, in order to overcome multilingualism and alternative names issues. Some of these vocabularies were already available and in use by the community: in this case, our contribution consists in gathering them, converting in SKOS and aligning them. As a result, we collected, implemented and published 23 controlled vocabularies belonging to 18 different categories (musical keys, types of derivation, modes, thematic catalogues, functions, musical genres, medium of performances, etc.) [20]. The vocabularies are all available in the DOREMUS triple store server via its public SPARQL endpoint.

The categories of genres and medium of performances contain each 6 different vocabularies, including well-established reference thesauri, as well as institution-specific lists. The vocabularies of these two categories have been aligned using a thesaurus matching tool [5].

The use of vocabularies opens up different possibilities, like the definition of labels in different languages or of alternate lemmata in the same language (i.e. the French terms “ut majeur” and “do majeur” which both refer to the key of *C major*). Different kinds of relationships between terms are defined, such as hierarchies (for example, “violin” is a narrower concept with respect to the “string” family). Previous research demonstrated how an RDF structure helps reasoning engines to discover links between different levels in the hierarchy of instruments [16].

4 Producing and interacting with a classical music knowledge graph

This section presents the development and publication of the DOREMUS classical music Knowledge Graph and various applications that make use of this graph: the exploratory search engine Overture, a knowledge-based recommender system, and a chatbot for query the database through natural language.

4.1 Building the DOREMUS Knowledge Graph

The DOREMUS ontology can be instantiated to populate a large Knowledge Base about classical music, containing data about artists, works, performances, scores and recordings. The graph gives access to fine-grained metadata coming from the music archives of important French cultural institutions, namely Radio France, the French National Library (BnF), and the Philharmonie de Paris.

These metadata were originally coming in several different data structures, such as XML serialisations and the librarian-specialised format MARC, that had to be converted in RDF according to the DOREMUS model.

The conversion pipeline relies on explicit expert-defined transfer rules (or mappings) that indicate where in the source file to look for what kind of information, providing the corresponding property path in the model as well as useful examples that illustrate each transfer rule, as shown in Figure 5. The role of these rules goes beyond being a simple documentation for the source records, embedding also information on some librarian practices in the formalisation of the content (format of dates, agreements on the syntax of textual fields, default values if the information is absent).

UNIT OF INFORMATION	F22 Expression: Opus Number
PATH	F22 Self-Contained Expression U17 has opus statement M2 Opus Statement [U42 has opus number M12 Opus Number] + [U43 has opus subnumber M13 Opus Subnumber]
INTERMARC BNF	TUM : 144 $p, chain of digits TUM : 144 $p, chain of digits before the comma
TRANSFER RULE	Remove the abbreviation “Op.” before the number
EXAMPLE	144 $pOp. 352 --> M12 = 352 144 $pOp. 27, no 2 --> M12 = 27, M13 =2

Fig. 5. Example of mapping rules describing the opus number and sub-number of a work

The source record files are read and interpreted according to the mapping rules, creating the RDF graph representing the record content (e.g. a musical work, a performance, etc.). Then, a parsing of the text notes is performed in order to extract more structured data from the text. This amounts to do a knowledge-aware parsing, since we search in the string the information we want to instantiate from the model (i.e. the instrument from the casting notes, or the date and the publisher from the first publication note). The parsing is realised making use of empirically-defined regular expression, as well as comparison of the text tokens with vocabularies, among which the music vocabularies (Section 3.3), GeoNames for locations and places [38], and the ISNI dataset[9] for persons.

Finally, a *string2uri* component, inspired by a module from the Datalift platform [34], performs an automatic mapping of string literals to URIs coming from controlled vocabularies. All variants for a concept label are considered in order to deal with potential differences in naming terms. As additional feature,

[9] http://www.isni.org/

this component is able to recognise and correct some noise that is present in the source file: this is the case of musical keys declared as genre, or fields for the opus number that contain actually a catalogue number and vice-versa. These cases and other typos and mistakes have been identified thanks to the conversion process and the visualisation of the converted data, supporting the partner institution in they work of updating and correcting constantly their data.

Currently, the DOREMUS Knowledge Graph includes more than 90 million triples, which describe over 18 million distinct entities. The classes and properties used come mostly from the DOREMUS ontology, FRBRoo and CIDOC-CRM. More details about the development of the DOREMUS graph is described in [1].

4.2 Interacting with the DOREMUS knowledge graph

User access to information is nowadays declined in several media which enable different and complementary approaches.

Overture*: an exploratory search engine.* Knowledge discovery is often entrusted to exploratory search engines. Instead of obtaining a precise result, the goal of exploratory search is learning something about a more or less vague topic, with a serendipitous attitude that pushes into continuing the search [27].

We developed an exploratory search engine for DOREMUS data, under the name of Overture (Ontology-driVen Exploration and Recommendation of mUsical REcords) [19]. The user is invited to explore the graph selecting one of the main concepts of the DOREMUS model – works, performances, scores, artists. In each of these sections, it is possible to perform advanced search using facets, making it possible to select precise subsets of data, like all sonatas that involve a clarinet and a piano. The hierarchical properties in controlled vocabularies allow the smart retrieval not only of entities matching exactly the chosen value (i.e. *strings*), but also any of its narrower concepts (i.e. *violin*, *cello*, etc.), taking into account also the interlinks between vocabularies. The multilingualism of the vocabularies enabled the application to be served in English and French.

Figure 6 depicts Beethoven's *Sonata for piano and cello n.1* as seen in Overture, as an example of a detailed page. Aside from the different versions of the title, the composer and a textual description, the page provides details on the information we have about the work. When these values come from a controlled vocabulary, a link is present in order to search for expressions that share the same value – e.g. the same genre – providing the user with a graph browsing experience. A timeline shows the most important events related to the work. Other performances and publications can be represented below and it is possible to click on them for accessing their detailed page.

Overture has largely been used for manually testing the content of the DOREMUS graph and the result of data conversion since its first development. The application is available at http://overture.doremus.org.

Recommending by graph similarity. On the right side in Figure 6, the UI presents similar items which the user may want to visit for continuing its exploration. We consider exploration and recommendation as two sides of the same

Fig. 6. OVERTURE: work detail

coin. With the first one, we let the user browse the datasets, discover connections on his own. Through recommendation, we remove this responsibility to the user with the purpose of presenting what he needs in a particular moment.

What do we suggest to a user listening to Beethoven? Similar works can share genres, instruments, composers and composition period, and so on. But how to define a similarity measure that take into account these concepts? Graph structures are particularly suitable for discovering connections between nodes. This is valid also for the DOREMUS knowledge graph, in which entities are linked through lower-level nodes.

We propose a solution based on graph embeddings generated at two different levels. The *feature embeddings* are computed at the level of the simple features (i.e. genre, instruments, etc.) directly on the knowledge graph, taking into account:

– the *graph of vocabularies*, which defines structural and semantic connections between entities, such as hierarchies, `owl:sameAs` links, properties in common, specific music properties. Given that this information has been redacted by human experts according to logic or historical reasons, it represents the involved concepts for what they are;
– the *graph of usage*, which includes all the usages of the vocabularies in the DOREMUS dataset. We considered musical works for the genre and the key, castings and performances for instruments, composition and performance

events for functions. This information represents the involved concept for how they occur in the reality of compositions and performances.

The *entity embeddings* represent the main entities (artists, works) and are the result of the combination of *feature embeddings*. In the example of works, the involved features are composition date, genre, key, etc.

Using graph embeddings reduces the similarity problem as the reverse of an Euclidean distance metrics which as been weighted prioritising the most homogeneous features in editorial playlists. The recommendation is then realised by selecting the most similar items [22]. An evaluation is performed by a group of music expert, revealing that the weights extracted with the analysis of editorial playlist improve the quality of the recommended items' ranking [18]. Future experiments will investigate the impact of this method as preliminary filter or fine-tuning module for state-of-art recommender systems.

DOREMUS Bot: ask classical music related questions in French and in English. In recent years, we have witnessed the rise of voice-based AI, which massively reached our homes and brought the knowledge available on the Web via voice calls. One the challenges related to this trending technology involves the design and development of smart conversational agents or chatbots [15], able to mimic the human conversation flow. Following the trend, we exposed part of the DOREMUS knowledge through a chatbot application, available at https://chatbot.doremus.org.

The bot is able to successfully recognise and answer different categories of questions (intents), allowing the user to retrieve information about works with given characteristics, artists searched by name or genre, and next performances of a composer's work in a given city.

Beyond being a way to publicly expose the DOREMUS data, the development of the chatbot allowed us to further validate the relevance of DOREMUS controlled vocabularies. The application makes strongly use of multi-language dictionaries of genre, instruments, and musicians, which are directly extracted from the DOREMUS endpoint. Their presence allowed to expose the chatbot in English and French and take into account all the different synonyms. In addition, a spell-checking module has been developed for detecting and correcting misspelled elements, acting in the context of each dictionary.

5 Conclusion and Future Work

In this chapter, we have presented the DOREMUS ontology, an extension of FRBRoo for the music domain. The model has a very rich expressiveness that allows, for instance, to describe different kinds of contributors (not only authors or performers), to detail the casting of a composition (with number, roles, notes for each instrument/voice), to specify performers at level of single performance inside a whole concert. This statement is supported by a series of specific competency questions which get an answer by querying the model. The DOREMUS ontology is used to populate the DOREMUS Knowledge Graph, containing information about classical music in RDF. The access to the data from the final user is ensured by an exploratory search engine, a recommender system, and a chatbot that exploit the strength of the Knowledge Graph (multilingualism, structured vocabularies, completeness).

On the other hand, the DOREMUS model is quite complex and hard to adopt if we look at the levels of distribution of information: mentioning an example, from an Expression, one has to pass through Event and Activity to reach a composer (see Figure 4). This complexity is indeed the heritage of both FRBR and CIDOC-CRM. The DOREMUS ontology defines 83 classes and 165 properties, which should be added to the 48 classes and 74 properties introduced by FRBRoo on top of the 84 classes and 161 properties of CIDOC-CRM[10], for a total number much higher than the one of the Music Ontology (54 classes and 153 properties). According to the *Linked Art* community[11], for 90% of times, only 10% of the full complexity of CIDOC-CRM is used. The dualism Work - Expression increases the number of required entities and triples for describing each part of the music information, often not really carrying significant extra information[12]. It is interesting to note that other FRBR-inspired models – like Music Ontology – prefer to skip this difference and propose a unique entity `MusicalWork` which puts together the two elements.

Another negative heritage of the extended models is the name convention for classes and properties, which foresees a succession of an uppercase letter, a number and the name of the class or the property, the latter always expressed as a verb. For this reason, DOREMUS recommends names like `U54 is performed expression of` in place of the shorter `performance of` of Music Ontology, or `E7 Activity` rather then `Activity` as in *Linked.art*. This impacts on query readability and speed issues. Some properties like `R17 created`, `R18 created`, `R21 created`, `P94 has created` consist in duplicates of the same action applied to different domains or ranges, making the model more error prone. Finally, the absence of a specific Music Work class[13] turned out to be an impacting problem, making hard to distinguish music pieces from other kind of works like text used in the lyrics, artistic objects used in scenes, etc.

A set of elements that are strictly connected to a librarian and cataloguer vision of the music object are included both in FRBRoo (e.g. `F40 Identifier Assignment`) and DOREMUS (e.g. `U172 has statement of responsibility relating to title`), introduced by the need for tracking the original source of specific statements. The results is mixture of metadata – the ones describing the music and the ones describing the metadata of the music – and in general could make the model be considered too librarian-specific. Further work could overcome this mixture by experimenting new Semantic Web approaches like

[10] These numbers do not include inverse properties

[11] https://linked.art/

[12] A common example is for entities of type *F14 Individual Work*, which quite often are just linked to the Expression, the Expression Creation and the provenance information, like in http://data.doremus.org/work/7259a748-6dd2-3e3d-b9de-7617d0a2b794.

[13] In the dataset, music work are `F14 Individual Work` with a type 'musical work'.

RDF* and SPARQL* [13]. These approaches enable the annotation of a RDF triple, using it in turn as subject or object of an RDF predicate. In this way, an additional layer of information is created which keeps separated the two levels of information.

All these reasons may potentially hamper the adoption of the DOREMUS model by a large public. The simplification of the ontology – for which a first attempt has been performed using the Schema.org vocabulary [21] – is therefore crucial and requires further work.

References

1. Achichi, M., Lisena, P., Todorov, K., Troncy, R., Delahousse, J.: DOREMUS: A Graph of Linked Musical Works. In: 17^{th} International Semantic Web Conference (ISWC). Monterey, CA, USA (2018)
2. Adamou, A., Brown, S., Barlow, H., Allocca, C., D'aquin, M.: Crowdsourcing Linked Data on Listening Experiences through Reuse and Enhancement of Library Data. Int. J. Digit. Libr. **20**(1), 61–79 (Mar 2019). https://doi.org/10.1007/s00799-018-0235-0
3. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: DBpedia: A Nucleus for a Web of Open Data. In: 6^{th} International Semantic Web Conference (ISWC) and 2^{nd} Asian Semantic Web Conference (ASWC). pp. 722–735. Springer-Verlag, Busan, Korea (2007), http://dl.acm.org/citation.cfm?id=1785162.1785216
4. Bainbridge, D., Hu, X., Downie, J.S.: A Musical Progression with Greenstone: How Music Content Analysis and Linked Data is Helping Redefine the Boundaries to a Music Digital Library. In: 1^{st} International Workshop on Digital Libraries for Musicology (DLfM). pp. 1–8. ACM (2014)
5. Bellahsene, Z., Emonet, V., Ngo, D., Todorov, K.: YAM++ Online: A Web Platform for Ontology and Thesaurus Matching and Mapping Validation. In: 14^{th} Extended Semantic Web Conference (ESWC), Poster & Demo Track. pp. 137–142. Springer, Portoroz, Slovenia (2017)
6. Cherfi, S.S.s., Guillotel, C., Hamdi, F., Rigaux, P., Travers, N.: Ontology-based annotation of music scores. In: Knowledge Capture Conference (K-CAP). pp. 10:1–10:4. ACM, Austin, TX, USA (2017). https://doi.org/10.1145/3148011.3148038
7. Choffé, P., Leresche, F.: DOREMUS: Connecting Sources, Enriching Catalogues and User Experience. In: 24^{th} IFLA World Library and Information Congress. Columbus, OH, USA (2016)
8. Daquino, M., Daga, E., d'Aquin, M., Gangemi, A., Holland, S., Laney, R., Meroño-Peñuela, A., Mulholland, P.: Characterizing the Landscape of Musical Data on the Web: State of the Art and Challenges. In: 2^{nd} Workshop on Humanities in the Semantic Web (WHiSe). Vienna, Austria (2017)
9. Doerr, M.: The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata. AI magazine **24**(3), 75–75 (2003)
10. Doerr, M., Bekiari, C., LeBoeuf, P.: FRBRoo: a conceptual model for performing arts. In: CIDOC Annual Conference. pp. 6–18 (2008)
11. Estermann, B.: "OpenGLAM" in Practice–How Heritage Institutions Appropriate the Notion of Openness. In: 20^{th} International Research Society for Public Management Conference (IRSPM). pp. 13–15. Hong Kong, China (2016)

12. Gomez-Perez, A., Fernández-López, M., Corcho, O.: Ontological Engineering: with examples from the areas of Knowledge Management, e-Commerce and the Semantic Web. Springer Science & Business Media (2006)
13. Hartig, O.: The RDF* and SPARQL* Approach to Annotate Statements in RDF and to Reconcile RDF and Property Graphs. In: W3C Workshop on Web Standardization for Graph Data. Berlin, Germany (2019)
14. Haslhofer, B., Isaac, A.: data.europeana.eu: The Europeana Linked Open Data Pilot. In: International Conference on Dublin Core and Metadata Applications (DC). vol. 0, pp. 94–104. Dublin Core Metadata Initiative, The Hague, Netherlands (2011)
15. Klopfenstein, L.C., Delpriori, S., Malatini, S., Bogliolo, A.: The Rise of Bots: A Survey of Conversational Interfaces, Patterns, and Paradigms. In: ACM SIGCHI Conference on Designing Interactive Systems (DIS). pp. 555–565. ACM, Edinburgh, United Kingdom (2017). https://doi.org/10.1145/3064663.3064672
16. Kolozali, S., Barthet, M., Fazekas, G., Sandler, M.B.: Knowledge Representation Issues in Musical Instrument Ontology Design. In: 12^{th} International Society for Music Information Retrieval Conference (ISMIR). pp. 465–470 (2011)
17. Lasar, M.: Digging into Pandora's Music Genome with musicologist Nolan Gasser. Ars Technica (2011), https://arstechnica.com/tech-policy/2011/01/digging-into-pandoras-music-genome-with-musicologist-nolan-gasser/
18. Lisena, P.: Knowledge-based music recommendation: Models, algorithms and exploratory search. Ph.D. thesis, EURECOM, Sophia Antipolis, France (2019)
19. Lisena, P., Achichi, M., Fernandez, E., Todorov, K., Troncy, R.: Exploring Linked Classical Music Catalogs with OVERTURE. In: 15^{th} International Semantic Web Conference (ISWC). Kobe, Japan (2016)
20. Lisena, P., Todorov, K., Cecconi, C., Leresche, F., Canno, I., Puyrenier, F., Voisin, M., Troncy, R.: Controlled Vocabularies for Music Metadata. In: 19^{th} International Society for Music Information Retrieval Conference (ISMIR). Paris, France (2018)
21. Lisena, P., Troncy, R.: DOREMUS to Schema.org: Mapping a Complex Vocabulary to a Simpler One. In: 20^{th} International Conference on Knowledge Engineering and Knowledge Management (EKAW) (2016), posters & Demos Track
22. Lisena, P., Troncy, R.: Combining Music Specific Embeddings for Computing Artist Similarity. In: 18^{th} International Society for Music Information Retrieval Conference (ISMIR), Late-Breaking Demo Track. Suzhou, China (2017)
23. Meroño-Peñuela, A., Daquino, M., Daga, E.: A Large-Scale Semantic Library of MIDI Linked Data. In: 5^{th} International Conference on Digital Libraries for Musicology (DLfM). Paris, France (2018)
24. Meroño-Peñuela, A., Hoekstra, R., Gangemi, A., Bloem, P., de Valk, R., Stringer, B., Janssen, B., de Boer, V., Allik, A., Schlobach, S., et al.: The MIDI Linked Data Cloud. In: 16^{th} International Semantic Web Conference (ISWC). pp. 156–164. Springer, Vienna, Austria (2017)
25. Nurmikko-Fuller, T., Bangert, D., Abdul-Rahman, A.: All the Things You Are: Accessing An Enriched Musicological Prosopography Through JazzCats. In: Digital Humanities. Montreal, Canada (2017)
26. Nurmikko-Fuller, T., Dix, A., Weigl, D.M., Page, K.R.: In Collaboration with In Concert: Reflecting a Digital Library As Linked Data for Performance Ephemera. In: 3^{rd} International Workshop on Digital Libraries for Musicology (DLfM). pp. 17–24. ACM, New York, NY, USA (2016)
27. Palagi, E., Gandon, F., Giboin, A., Troncy, R.: A Survey of Definitions and Models of Exploratory Search. In: ACM Workshop on Exploratory Search and Interactive Data Analytics (ESIDA). pp. 3–8. ACM, Limassol, Cyprus (2017)

28. Raimond, Y., Abdallah, S.: The Event Ontology. Tech. rep., Queen Mary University of London (2007), http://motools.sourceforge.net/event
29. Raimond, Y., Abdallah, S.: The Timeline Ontology. Tech. rep., Queen Mary University of London (2007), http://purl.org/NET/c4dm/timeline.owl
30. Raimond, Y., Abdallah, S.A., Sandler, M.B., Giasson, F.: The Music Ontology. In: 15^{th} International Conference on Music Information Retrieval (ISMIR). pp. 417–422. Vienna, Austria (2007)
31. Raimond, Y., Sandler, M.B.: A Web of Musical Information. In: 9^{th} International Conference of the Society for Music Information Retrieval (ISMIR). pp. 263–268 (2008)
32. Rashid, S.M., De Roure, D., McGuinness, D.L.: A Music Theory Ontology. In: 1^{st} International Workshop on Semantic Applications for Audio and Music (SAAM). pp. 6–14. ACM, Monterey, CA, USA (2018)
33. Ren, Y., Parvizi, A., Mellish, C., Pan, J.Z., van Deemter, K., Stevens, R.: Towards competency question-driven ontology authoring. In: Presutti, V., d'Amato, C., Gandon, F., d'Aquin, M., Staab, S., Tordai, A. (eds.) The Semantic Web: Trends and Challenges. pp. 752–767. Springer International Publishing, Cham (2014)
34. Scharffe, F., Atemezing, G., Troncy, R., Gandon, F., Villata, S., Bucher, B., Hamdi, F., Bihanic, L., Képéklian, G., Cotton, F., et al.: Enabling linked-data publication with the datalift platform. In: AAAI workshop on semantic cities (2012)
35. Schedl, M.: Towards Personalizing Classical Music Recommendations. In: 2^{nd} International Workshop on Social Media Retrieval and Analysis (SoMeRA). Atlantic City, NJ, USA (2015)
36. Thomas, B., Bermes, E., Coyle, K., Dunsire, G., Isaac, A., Murray, P., Panzer, M., Schneider, J., Singer, R., Summers, E., Waites, W., Young, J., Zeng, M.: W3C Library Linked Data Incubator Final Report. Tech. rep., World Wide Web Consortium (W3C) (2011)
37. Vila-Suero, D., Gómez-Pérez, A.: datos.bne.es and MARiMbA: an insight into library linked data. Library Hi Tech **31**(4), 575–601 (2013). https://doi.org/10.1108/LHT-03-2013-0031
38. Wick, M., Vatant, B.: The GeoNames geographical database (2012), http://geonames. org