

SEMANTIC INTEGRATION OF MIR DATASETS WITH THE POLIFONIA ONTOLOGY NETWORK

Valentina Anita Carriero¹ Fiorela Ciroku¹ Jacopo de Berardinis²
Delfina Sol Martinez Pandiani¹ Albert Meroño-Peñuela² Andrea Poltronieri¹ Valentina Presutti¹

¹ University of Bologna, Italy

² King's College London, United Kingdom

andrea.poltronieri2@unibo.it

ABSTRACT

Integration between different data formats, and between data belonging to different collections, is an ongoing challenge in the MIR field. Semantic Web tools have proved to be promising resources for making different types of music information interoperable. However, the use of these technologies has so far been limited and scattered in the field. To address this, the Polifonia project¹ is developing an ontological ecosystem that can cover a wide variety of musical aspects (musical features, instruments, emotions, performances). In this paper, we present the Polifonia Ontology Network, an ecosystem that enables and fosters the transition towards Semantic MIR.

1. INTRODUCTION

During the last 20 years, the field of Music Information Retrieval (MIR) has seen the introduction of an increasing number of music datasets, enabling researchers to train and evaluate algorithms for several tasks, from chord recognition and beat detection, to source separation and mood detection. However, the availability of audio data is still limited, and two overlooked issues remain: (i) music datasets are commonly provided as independent and isolated collections, with little or no alignment at the metadata and annotation level; (ii) even when tracks/compositions are coupled with universal identifiers (e.g. MusicBrainz IDs, ISRC), there is no direct way to access and link heterogeneous music-related data from online databases, such as Wikipedia, Genius [1], and Songfacts [2]. The disconnect among music datasets jeopardises their potential integration, and hence their extension and the combination of annotations of different kinds. Simultaneously, the low level of linkage with other databases discourages multi-modal

research in the field, where the availability of heterogeneous music-related data (text, images, locations, etc.) is an essential asset. Therefore, MIR practitioners interested in multi-modal and/or multi-task research across existent music collections are commonly left to implement complex data collection and integration pipelines [3].

Semantic Web (SW) technologies [4] allow many of these limitations to be overcome, as they provide a common framework for data to be shared and reused across application, enterprise, and community boundaries. In the Resource Description Framework [5], entities can be described through triples of the form *subject-predicate-object*, and can have unique IDs – allowing to connect them to other entities through additional triples and form a Knowledge Graph (KG). Because not only documents but any entity of interest – such as musical resources – can be annotated, SW technologies greatly facilitate effective data access and integration, resource discovery, semantic reasoning and knowledge extraction, while also promoting interoperability among resources, information models, data providers and consumers. As such, they have been adopted by both large corporations, such as Oracle, IBM, Google²), and many domains dealing with vast quantities of data gathered with different encoding formats, such as health care, life sciences, and cultural heritage³. A wealth of recent research is dedicated to the adoption of SW knowledge bases as references for enhancing Information Extraction tasks [6].

In light of this, SW technologies and principles are ideal for MIR, as they could easily address the disconnection and the low level of linkage of music collections. Nonetheless, despite the numerous music ontologies to date, high fragmentation and poor maintenance of these contributions are currently hindering the transition to Semantic MIR (see Section 2). To fill this gap, we are developing an ontological ecosystem providing a standard, flexible and expressive schema to represent and describe/annotate heterogeneous musical data of different formats, genres, and provenance. By leveraging the different sets of skills and expertise of music scholars in Polifonia, we aim at unifying views and identifying requirements across different disciplines to guide the ontology design activities, while following state of the art methodologies for data engineering.

¹ <https://polifonia-project.eu>



© V.A. Carriero, F. Ciroku, J. de Berardinis, D.S. Martinez Pandiani, A. Meroño-Peñuela, A. Poltronieri, and V. Presutti. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** V.A. Carriero, F. Ciroku, J. de Berardinis, D.S. Martinez Pandiani, A. Meroño-Peñuela, A. Poltronieri, and V. Presutti, “Semantic Integration of MIR Datasets with the Polifonia Ontology Network”, in *Extended Abstracts for the Late-Breaking Demo Session of the 22nd Int. Society for Music Information Retrieval Conf.*, Online, 2021.

² <https://schema.org>

³ <https://w3.org/2001/sw/sweo/public/UseCases/>

2. MUSIC ONTOLOGIES

In the last two decades several ontologies have been developed for diverse music-related applications, dealing with both symbolic notations and audio signals at different levels of specificity. Some ontologies have been designed for describing high-level music-related information, such as the The Music Ontology [7] and the DOREMUS Ontology [8]. Other ontologies describe musical notation, both from scores and symbolic representations. For example, the MIDI Linked Data Cloud [9] proposes the interconnection of symbolic music descriptions encoded in MIDI format, and the CHARM ontology [10] aims to describe musical structures based on the CHARM specifications. The Music Theory Ontology (MTO) [11] describes theoretical concepts related to a music composition, while the Music Score Ontology [12] represents similar concepts with a focus on music sheet notation. Finally, the Music Notation Ontology [13] focuses on the core “semantic” information present in a score. Other ontologies aim to describe specific aspects of the musical domain, such as the Chord Ontology, the Tonality Ontology, the Temperament Ontology [14], and the Segment Ontology [15]. The Audio Features Ontology [16], the Studio Ontology [17], and the Audio Effects Ontology [18] describe audio signals and production procedures. Others have also been used to model listening habits and music tastes [19], music-induced emotions [20], and to describe musical similarities [21].

The focus of these ontologies is generally specific, covering only particular aspects of musical content. However, music consists of a dense connected network of heterogeneous elements that concert with each other. Furthermore, many of these ontologies were developed as stand-alone projects, with little or no alignment to other relevant ontologies within the same domain.

3. THE POLIFONIA ONTOLOGY NETWORK

Besides the expressive and modular design, allowing to represent a wide range of music-related concepts and relations, the Polifonia ontology network (ON) also brings two key assets desirable for the whole ecosystem. First, music objects are described in a way to be connected with relevant knowledge, such as their links to tangible objects, cultural and historical contexts, stories about them, as well as related facts expressed in different styles, disciplines, languages. This will enable the development of an ecosystem of computational methods and tools supporting discovery, extraction, encoding, interlinking, classification, exploration of, and access to, musical heritage knowledge on the Web. Second, the development of Polifonia ON⁴ is driven by eXtreme Design (XD) [22] – an agile ontology engineering methodology, and makes extensive use of ontology design patterns (ODPs) – small ontologies that work as reusable templates for recurrent modelling problems. For example, the common modelling problem in the music domain “Where was a musical composition performed?” could be addressed by an ODP modelling the

relation between compositions, performances, and places.

XD is executed through multiple iterations of a set of collaborative steps: (i) requirements collection; (ii) ontology design and KG generation; (iii) ontology testing and integration. A story-based approach guides the collection of the project’s requirements. A *story* is a framework for *customers* to describe their needs, and is composed of 4 sections: the *persona*, a description of a typical user; the *goal* that she needs to address; the *scenario*, describing how she will address the goal(s); the *competency questions* (CQs), that translate the needs of the persona into requirements. The ontology modelling starts iteratively from the CQs, and is based on the reuse of ODPs and entities extracted from existing ontologies as templates. For instance, if an existing ontology already addresses one of the CQs, alignment axioms⁵ make it evident which part(s) have been reused in Polifonia ON and support interoperability between ontologies. This makes Polifonia ON and KG interoperable and reusable beyond the Polifonia project.

The identification of the 11 ontology modules emerged from the investigation, and thematic classification, of all the CQs from the 19 stories recorded so far. The *Full* module is the *entry point* of the whole network. The *Core* module models general-purpose concepts and relationships (e.g. place, time). The *Musical Performance* and *Musical Composition* modules represent musical performances and events, and musical compositions respectively. The *Musical Feature* module provides a model to describe a musical object in regard to the musical properties that can be objectively attributed or subjectively identified from it. The *Source* module represents sources of music-related information. The *Instrument* module proposes a taxonomy of instruments and their technical properties, whereas *Bell* focuses on bells. The *Music emotion* module provides a model to describe emotions both perceived and induced by a composition wrt the musical features. A module named *Comparative Measure* defines ODPs describing observations and measures applied to comparative analysis, such as similarities between musical pieces. Lastly, the *Metadata* module aims at supporting the representation of metadata about musical resources. The Polifonia ON is being populated by data from various existing datasets.

4. CONCLUSIONS AND FUTURE WORK

In this paper, we addressed the recurring problem of integration and linkage of MIR datasets. By means of SW technologies and best practices, we are laying the foundations and providing the infrastructure for Semantic MIR. Not only will this support computational music analysis, but the KGs resulting from the interconnection of MIR datasets and their integration with other sources can be explored through symbolic reasoning to derive novel musical knowledge and test musicological hypotheses. We are currently developing a first version of the Polifonia ontology network, and future work will iteratively extend the ecosystem as a result of continuous expert feedback.

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

⁵ E.g. an alignment axiom could assert that the concept of *Person* in Polifonia has the same semantics as the concept of *Person* in FOAF.

Acknowledgements. The Polifonia project has received funding from the EU Horizon 2020 research and innovation programme under grant agreement No 101004746.

5. REFERENCES

- [1] “Genius website,” <http://www.genius.com>, accessed: 2021-10-30.
- [2] “Songfacts website,” <http://www.songfacts.com>, accessed: 2021-10-30.
- [3] I. A. P. Santana, F. Pinhelli, J. Donini, L. Catharin, R. B. Mangolin, V. D. Feltrim, M. A. Domingues *et al.*, “Music4all: A new music database and its applications,” in *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*. IEEE, 2020.
- [4] T. Berners-Lee, J. Hendler, and O. Lassila, “The semantic web,” *Scientific american*, vol. 284, no. 5, 2001.
- [5] O. Lassila, R. R. Swick *et al.*, “Resource description framework (rdf) model and syntax specification,” 1998.
- [6] J. L. Martinez-Rodriguez, A. Hogan, and I. Lopez-Arevalo, “Information extraction meets the semantic web: a survey,” *Semantic Web*, vol. 11, no. 2, 2020.
- [7] Y. Raimond, S. Abdallah, M. Sandler, and F. Giasson, “The music ontology,” in *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007)*, Vienna, Austria, Sep. 2007.
- [8] P. Lisena and R. Troncy, “Doing reusable musical data (DOREMUS),” in *Proceedings of Workshops and Tutorials of the 9th International Conference on Knowledge Capture (K-CAP2017)*, Austin, Texas, USA, December 4th, 2017, ser. CEUR Workshop Proceedings, vol. 2065. CEUR-WS.org, 2017.
- [9] A. Meroño-Peñuela, R. Hoekstra, A. Gangemi, P. Bloem, R. de Valk, B. Stringer, B. Janssen, V. de Boer, A. Allik, S. Schlobach, and K. Page, “The midi linked data cloud,” in *The Semantic Web – ISWC 2017*. Cham: Springer International Publishing, 2017.
- [10] N. Harley and G. Wiggins, “An ontology for abstract, hierarchical music representation,” in *Demo at the 16th International Society for Music Information Retrieval Conference (ISMIR 2015)*, Malaga, Spain, 2015.
- [11] S. M. Rashid, D. De Roure, and D. L. McGuinness, “A music theory ontology,” in *Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music*, ser. SAAM ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 6–14.
- [12] J. Jones, D. de Siqueira Braga, K. Tertuliano, and T. Kauppinen, “Musicowl: The music score ontology,” in *Proceedings of the International Conference on Web Intelligence*, ser. WI ’17. New York, NY, USA: Association for Computing Machinery, 2017.
- [13] S. S.-s. Cherfi, C. Guillotel, F. Hamdi, P. Rigaux, and N. Travers, “Ontology-based annotation of music scores,” in *Proceedings of the Knowledge Capture Conference*, ser. K-CAP 2017. New York, NY, USA: Association for Computing Machinery, 2017.
- [14] G. Fazekas, Y. Raimond, K. Jacobson, and M. Sandler, “An overview of semantic web activities in the omras2 project,” *Journal of New Music Research*, vol. 39, 12 2010.
- [15] B. Fields, K. R. Page, D. D. Roure, and T. Crawford, “The segment ontology: Bridging music-generic and domain-specific,” in *Proceedings of the 2011 IEEE International Conference on Multimedia and Expo, ICME 2011, 11-15 July, 2011, Barcelona, Catalonia, Spain*. IEEE Computer Society, 2011.
- [16] A. Allik, G. Fazekas, and M. B. Sandler, “An ontology for audio features,” in *Proceedings of the 17th International Society for Music Information Retrieval Conference, ISMIR 2016, New York City, United States, August 7-11, 2016*, 2016.
- [17] G. Fazekas and M. B. Sandler, “The studio ontology framework,” in *Proceedings of the 12th International Society for Music Information Retrieval Conference, ISMIR 2011, Miami, Florida, USA, October 24-28, 2011*. University of Miami, 2011.
- [18] T. Wilmering, G. Fazekas, and M. B. Sandler, “The audio effects ontology,” in *Proceedings of the 14th International Society for Music Information Retrieval Conference, ISMIR 2013, Curitiba, Brazil, November 4-8, 2013*, 2013.
- [19] S. Rho, S. Song, E. Hwang, and M. Kim, “Comus: Ontological and rule-based reasoning for music recommendation system,” in *Advances in Knowledge Discovery and Data Mining*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009.
- [20] H. H. Kim, “A semantically enhanced tag-based music recommendation using emotion ontology,” in *Intelligent Information and Database Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.
- [21] K. Jacobson, Y. Raimond, and M. B. Sandler, “An ecosystem for transparent music similarity in an open world,” in *Proceedings of the 10th International Society for Music Information Retrieval Conference, ISMIR 2009, Kobe International Conference Center, Kobe, Japan, October 26-30, 2009*. International Society for Music Information Retrieval, 2009.
- [22] E. Blomqvist, K. Hammar, and V. Presutti, “Engineering ontologies with patterns-the extreme design methodology.” *Ontology Engineering with Ontology Design Patterns*, no. 25, 2016.