Linking MIDI Data for multidisciplinary and FAIR AI on Music

Albert Meroño-Peñuela, Vrije Universiteit Amsterdam

More and more researchers from a large variety of fields are using modern AI techniques to understand, describe, and generate musical content and metadata. In the last years we have seen a blooming diversity of algorithms, mainly learned automatically from data using deep learning architectures, to address various tasks in Music Information Retrieval such as instrument separation, automatic transcription, and chord detection. Some of these architectures use the models they learn to generate genuinely new, but still data-inspired, songs and tunes that start posing questions about machine vs human musical creativity.[1] In more musicological and humanistic settings, fundamental questions and hypotheses are typically multidisciplinary in nature, normalling dealing with aspects of culture, economy, and musical traits and characteristics. The common requirement for methods in all these fields is the same: large, carefully curated, and explicitly and semantically described *datasets*.

Generating these datasets in such a way that they can be reused among different asks in those various fields is challenging. There exist several reasons for this: (1) existing musical data is distributed and scattered all over the Web; (2) meaningful connections between these data are missing (e.g. a Spotify song URI and its corresponding URI entry in MusicBrainz[2]); (3) equivalent representations of musical knowledge, like symbolic music scores, are encoded in various formats (MIDI, MusicXML, MEI, LillyPond, ABC, etc.) that are hardly interoperable.

Solving these issues at scale for arbitrary data (i.e. not just musical information) has been one the driving goals of 20 years of Semantic Web research [6]. This research has reached now maturity and delivered technologies to make Knowledge Graphs possible; these Knowledge Graphs, when deployed, integrate, interlink, and make datasets interoperable on the Web. In this poster, we will reflect on current work towards the construction of Music Knowledge Graphs that integrate symbolic musical content in the MIDI format and metadata databases that describe and enrich that content. In the first of such works, we proposed the *midi2rdf* algorithm [1];[3] a program that can lossless convert any individual MIDI file into the interoperable Resource Description Framework (RDF) format through a low-level ontology. Following up on this, we used community-based knowledge from GitHub to gather a comprehensive list of sites containing MIDI files on the Web (around 500K);[4] and run the *midi2rdf* algorithm on them to create the MIDI Linked Data Cloud[5], a MIDI Knowledge Graph of 10B semantic statements [2]. The first creative use of such integrated musical knowledge space was to enable large scale mashups (mixes of different tracks of various songs), in which the compatibility of different candidate tracks for the mix (e.g. their rhythmic similarity) is automatically determined by a SPARQL pattern [3] by the so-called "SPARQL-DJ"[6].

Building on recent deep learning architectures, the Text2MIDI Variational Auto Encoder [5] learns the relationship between words in lyrics and melodic motifs, generating melodies automatically for arbitrary input poems.[7]

Three challenging tasks remain open for the future. The first is to envisage algorithms for automatically linking MIDI musical content --or any other musical symbolic format-- to musical metadata databases (e.g. MusicBrainz[8]) at scale, in order to enable a richer discoverability and retrieval of related resources. First steps towards this goal are already made in the Semantic Web MIDI Tape [4] through amateur playback, MIDI similarity algorithms, and metadata propagation.[9] The second challenge is to expand the interoperability of Linked Data Musical dataspaces to other symbolic notation formats, like ABC, LillyPond, MusicXML or MEI, which remain hardly interoperable. The third challenge is to better understand the integration requirements from the various fields doing research in AI and Music in order to conceive appropriate strategies, data models, and access methods for Music Knowledge Graphs.

References

1. Albert Meroño-Peñuela, Rinke Hoekstra. "The Song Remains the Same: Lossless Conversion and Streaming of MIDI to RDF and Back". In: 13th Extended Semantic Web Conference (ESWC 2016), posters and demos track. May 29th — June 2nd, Heraklion, Crete, Greece (2016).
2. Albert Meroño-Peñuela, Rinke Hoekstra, Aldo Gangemi, Peter Bloem, Reinier de Valk, Bas Stringer, Berit Janssen, Victor de Boer, Alo Allik, Stefan Schlobach, Kevin Page. "The MIDI Linked Data Cloud". In: The Semantic Web – ISWC 2017, 16th International Semantic Web Conference. Lecture Notes in Computer Science, vol 10587, pp. 156-164 (2017).
3. Rick Meerwaldt, Albert Meroño-Peñuela, Stefan Schlobach. "Mixing Music as Linked Data: SPARQL-based MIDI Mashups". In: Proceedings of the 2nd Workshop on Humanities in the SEmantic web (WHiSe 2017). ISWC 2017, October 22nd, Vienna, Austria (2017).
4. Albert Meroño-Peñuela, Reinier de Valk, Enrico Daga, Marilena Daquino, Anna Kent-Muller. "The Semantic Web MIDI Tape: An Interface for Interlinking MIDI and Context Metadata". In: Workshop on Semantic Applications for Audio and Music, ISWC 2018. 9th October 2018, Monterey, California, USA (2018).
5. Roderick van der Weerdt. "Generating Music from Text: Mapping Embeddings to a VAE's Latent Space". MSc thesis, Universiteit van Amsterdam (2018).
6. Berners-Lee, Tim, James Hendler, and Ora Lassila. "The semantic web." Scientific american 284.5 (2001): 28-37.

---

[7] https://github.com/RoderickvanderWeerdt/text2MIDI
[8] https://musicbrainz.org/
[9] https://www.youtube.com/watch?v=xSfo3dwF1mY