

# FAIR Interconnection and Enrichment of Public-Domain Music Resources on the Web

DAVID M. WEIGL[1]

*University of Music and Performing Arts Vienna, Austria*

TIM CRAWFORD

*Goldsmiths, University of London, United Kingdom*

AGGELOS GKIOKAS

*Universitat Pompeu Fabra, Spain*

WERNER GOEBL

*University of Music and Performing Arts Vienna, Austria*

EMILIA GÓMEZ

*Universitat Pompeu Fabra, Spain*

NICOLÁS F. GUTIÉRREZ

*Universitat Pompeu Fabra, Spain*

CYNTHIA C. S. LIEM

*Delft University of Technology, The Netherlands*

PATRICIA SANTOS

*Universitat Pompeu Fabra, Spain*

**ABSTRACT:** Vast amounts of publicly licensed classical music resources are housed within many different repositories on the Web encompassing richly diverse facets of information—including bibliographical and biographical data, digitized images of music notation, music score encodings, audiovisual performance recordings, derived feature data, scholarly commentaries, and listener reactions. While these varied perspectives ought to contribute to greater holistic understanding of the music objects under consideration, in practice, such repositories are typically minimally connected. The TROMPA project aims to improve this situation by interconnecting and enriching public-domain music repositories. This is achieved, on the one hand, by the application of automated, cutting-edge Music Information Retrieval techniques, and on the other, by the development of contribution mechanisms enabling users to integrate their expertise. Information within established repositories is interrelated with data generated by the project within a data infrastructure whose design is guided by the FAIR principles of data management and stewardship: making music information *Findable, Accessible, Interoperable, and Reusable*. We provide an overview of challenges of *description, identification, representation, contribution, and reliability* toward applying the FAIR principles to music information, and outline TROMPA's implementational approach to overcoming these challenges. This approach applies a graph-based data infrastructure to interrelate information hosted in different repositories on the Web within a unifying data model (a 'knowledge graph'). Connections are generated across different representations of music content beyond the catalogue level, for instance connecting note elements within score encodings to corresponding moments in performance time-lines. Contributions of user data are supported via privacy-first mechanisms that retain control of such data with the contributing user. Provenance information is captured throughout, supporting reproducibility and re-use of the data both within and outside the context of the project.

Submitted 2020 April 30; accepted 2020 October 14.

Published 2021 December 10; <https://doi.org/10.18061/emr.v16i1.7643>

**KEYWORDS:** *Music information, data infrastructure, Linked Data, digital musicology*



CLASSICAL music is a key element of European cultural heritage: both a treasured legacy and an active, lived tradition that continues to be performed and enjoyed by many people today. Much of the classical repertoire has entered the public domain, and librarians, scholars, cultural heritage practitioners, and other enthusiasts have digitized and published massive amounts of scores and recordings on the Web.

These materials are a boon not only to music research. Historical (political, social, religious, economic, etc.), literary, artistic, archaeological or architectural studies also stand to benefit from the facilitated availability of such resources. Furthermore, studies of the actual performance of music—how the music was intended to sound in its original context and how this might be emulated in modern performance—need to combine this kind of broad historical and musicological research with the artistic endeavor of musical performance.

The evidence on which all this work is fundamentally based is to be found in the various types of documents stored in libraries: books, manuscripts, editions of source materials, articles in scholarly or other journals and newspapers, and all kinds of archival materials. Increasingly, this evidence is being made manifest in digital resources, becoming subject to analyses within the emergent, transdisciplinary field of ‘digital musicology’. In such digitization efforts, textual or music-score documents have typically been captured as PDFs or image-format files. Though convenient for storage, downloading, viewing, and printing by end-users such as musicians or scholars, such formats do not represent musical meaning in a machine-accessible form—a score image is a collection of pixels, not notes—presenting difficulties to automated processing of such materials by software algorithms.

Music performers and music enthusiasts also stand to benefit from the increasingly rich availability of classical music resources, in the form of music scores and vast collections of audio(-visual) performance recordings. However, the user generally has to manually discover correspondences between scores and recordings published on the Web, normally by reference to a work’s title. This process is time-consuming and requires significant expertise regarding materials and Web repositories. More dynamic ways of interacting with such sources, such as contributing own materials, commenting on specific sections of scores or performances, or engaging in comparative performance analyses are currently largely unsupported.

Though vast quantities of music information are available as open data, i.e., permissively licensed for public consumption and reuse, freely accessible to anyone, making these resources *Findable, Accessible, Interoperable*, and *Retrievable* in accordance with the FAIR principles for scientific data management and stewardship (Wilkinson et al., 2016) remains challenging. However, such practices must be followed if more advanced, dynamic, and holistic means of incorporating human insight and machine analyses across these diverse resources are to be supported.

The EU Horizon 2020-funded TROMPA project[2] (Towards Richer Online Music Public-domain Archives) is working towards this goal, combining Music Information Retrieval (MIR) technologies and crowd-sourcing approaches to publish, interlink, contextualize, and augment public-domain classical music resources. Here, we mean the term ‘classical music’ in broad, inclusive terms, mirroring the general public understanding. TROMPA incorporates large established online repositories’ collections, and provides services for discovery, enhancement, and contribution of musical scores, recordings, analyses, and interpretations. The tools and workflows developed within the project apply open, standard Web and MIR technologies to ensure reusable, reproducible, and re-interpretable access to the data produced. Within the project, we are also developing specialized user-facing Web applications (Weigl, Goebel et al., 2019) supporting interactions with these materials in music scholarship, music performance, and citizen science contexts.

In this paper, we describe the challenges faced by initiatives seeking to interlink and enrich public-domain classical music resources on the Web, and to provide useful, FAIR access to the data while respecting stakeholder rights. Critically, these include the rights of contributors and participants in accordance with the EU’s General Data Protection Regulation (GDPR; Voigt & von dem Bussche, 2017). TROMPA adopts the FAIR principles over other approaches to data management due to their increasingly widespread application to research data in recent years (van Reisen et al., 2020). Further, they are closely aligned with the Linked Data paradigm (Hasnain & Rebholz-Schuhman, 2018), in which structured information and assertions about resources stored on the Web are represented in interoperable, machine-interpretable form using uniform resource identifiers (URIs). This approach is core to TROMPA’s data infrastructure. Finally, the FAIR principles are explicitly incorporated within the funding requirements set out by the EU’s Horizon 2020 scheme.

The challenges we detail in this paper were chosen to best reflect experiences within TROMPA and its predecessor projects, which also involved large-scale semantic publication and processing of machine-

and user-generated music data and metadata (see also Lewis, Crawford, & Lewis, 2015; Liem, Gomez, & Schedl, 2015; Page, Bechhofer, Fazekas, Weigl, & Wilmering, 2017; Sandler, De Roure, Benford, & Page, 2019). In particular, we focus on challenges to *findability*, *accessibility*, *interoperability*, and *reusability* specifically arising in contexts where multimodal, multimedia information is gathered from numerous disparate sources (including human and machine agents), aggregated, interconnected, and enriched in machine-readable fashion, and made available on the Web. Further challenges that must be faced by any project providing large-scale, long-term publication of data (e.g., relating to database integrity and persistence) are explicitly declared out of scope. The challenges we detail here include:

- **Challenges of description:** Digitization projects undertaken by different institutions, projects, or individuals often vary fundamentally in the approach taken to describe, annotate, or catalog their data, even if the resources under consideration are very similar. For example, though two projects might concern the digitization of printed score sheets (even, perhaps, the same piece and edition), the vocabularies used to describe these resources and the data structures used to house the corresponding metadata may differ significantly. Depending on the individual publication, these metadata may not be exposed in their raw form but only through a custom application programming interface (API), a site-specific, non-generalizable set of functions; or they may not be exposed in machine-readable form at all, requiring web-scraping and text processing in order to recover (approximations of) this important information.
- **Challenges of identification:** Various repositories on the Web provide access to classical music in different modalities, but interlinking between such repositories tends to be limited. This is the problem of ‘data silos’: While different digital libraries, archives, and other Web-accessible repositories may provide access to complementary information, often describing different aspects of a shared underlying entity (e.g., different score editions, performance recordings, or scholarly commentaries of the same piece of music), unified viewing and analysis of these mutually informative facets is hindered by a lack of connectivity.
- **Challenges of representation:** A musical work may appear in many parallel forms (Liem, Müller, Eck, Tzanetakis, & Hanjalic, 2011; Wiggins, 2016): as historical and contemporary, professional and amateur recordings, that may either contain all movements or different fragments, but also in the form of symbolic notation, where an amateur player may ‘just want the notes’, but a scholar may want to compare a professional edition to the original manuscript or to parallel editions. Corresponding resources available in music Web repositories are typically manifested as digitized score images and audiovisual recordings which, while highly meaningful to music scholars, performers, and enthusiasts, are semantically undifferentiated from a machine perspective, highlighting the need for multifaceted, semantically-enriched representations of music information (Sandler et al., 2019).
- **Challenges of contribution:** Tasks pertaining to interlinking and enrichment of information resources in music repositories can be completed to some extent using automated MIR technologies. However, certain core activities require the application of scholarly insight, artistic interpretation, or subjective experience, and greatly benefit from the inclusion of user contributions. Handling such contributions poses significant challenges both from a data infrastructure and a data protection and intellectual property rights perspective, and in both aspects complicates the situation beyond the application of MIR algorithms.
- **Challenges of reliability:** Even where feature descriptions can be derived from music resources through the automated application of MIR processes, collections of such metadata do not always expose sufficiently detailed information pertaining to their provenance and processing context. However, such precision is required if empirical findings are to be reproduced, recontextualized, or re-interpreted outside of the immediate project scope to adequately support reliable empirical musicological analyses.

The TROMPA project is motivated by the pursuit of these challenges, working towards a vision of *Findable*, *Accessible*, *Interoperable*, and *Reusable* music information resources. In this vision, diverse collections of publicly licensed classical music resources available on the Web are mutually discoverable, and the multimodal information housed within and across these different repositories is made available for unified query, access, and analysis in order to provide complementary perspectives on a shared music repertoire. Beyond retrieval of such information, semantic Web technologies are incorporated to capture and expose musical meaning to machine as well as human agents; and, to allow contributions by music scholars,

performers, and enthusiasts to be welcomed in a scalable fashion, respecting and upholding their rights as contributors in order to better support the contemporary, living tradition of classical music.

In Sections 1 – 5, we will detail each of these challenges, mapping each to corresponding FAIR principles as listed on the website of the GO FAIR initiative that promotes their ongoing adoption, available online at <https://www.go-fair.org/fair-principles/>. In particular, we refer to each principle using the identifiers used on that website, printed in **bolded italics** – e.g., *F1* – “... a globally unique and persistent identifier”, or *RI.2* – “(Meta)data are associated with detailed provenance”. In Section 6, we then detail TROMPA’s approaches to addressing these challenges, before presenting concluding remarks in Section 7.

## 1. CHALLENGES OF DESCRIPTION

Many different organizations, projects, and people, including cultural heritage institutions, libraries, academic researchers, broadcasters, commercial companies, and enthusiast individuals, have published rich, multimodal repositories of music on the Web. Though sharing a common goal of providing access to music information, these endeavors have applied different approaches toward exposing this information for retrieval, complicating the task of providing unified views and querying over these disparate collections.

Notable repositories exposing public-domain music information include the International Music Score Library Project (IMSLP)[3], also known as the Petrucci Music Library, which houses digitized score images for more than 510,000 public-domain scores by more than 17,500 composers, alongside more than 60,000 publicly-licensed recordings and other resources. These are exposed to users through a website providing a query interface, and to software agents through a simple API[4] which provides machine-readable access to the collection but does not support searching. This makes it difficult to find specific resources, particularly as entities are not connected to external authority identifiers (Fields, Phippen, & Cohen, 2015).

The repository offered by Muziekweb[5], a large music library in the Netherlands, provides another instructive example. Its collection houses over 600,000 CDs and 300,000 LPs, described in conformance with international library standards guided by the Online Computer Library Center (OCLC), a global library cooperative[6]. The Muziekweb website offers convenient options for query and discovery of these descriptions. Their API[7] provides machine-readable access to the underlying data as well as complex searching along a number of descriptors. Interlinking with external authorities is currently limited, but it is a goal for medium-term development.

Another noteworthy example is MusicBrainz, a community-maintained ‘online music encyclopedia’ housing vast quantities of structured (machine-readable) music metadata. MusicBrainz’ venerable origins in the early days of the Web (Swartz, 2002), its focus on a fine-grained exposure of data through a comprehensive API service, and its inclusion of reciprocal linking with other repositories have positioned it as a widely used authority for music identification, used by many organizations including the British Broadcasting Corporation (BBC)[8].

The three above-mentioned repositories serve as multimodal examples of online, large-scale, widely used sources of public-domain music information. Of course, many more such repositories exist—and the approaches outlined in this paper (particularly in Section 6) aim to generalize to any such collection exposed (i.e., made addressable through URIs) on the Web. In TROMPA, these three repositories are chosen as pilot sources primarily providing access to information in three important modalities: scores, audio recordings, and catalogue metadata. In this section, these examples serve to illustrate challenges to the FAIR principles[9] of *findability* and *interoperability* in the music domain. Each repository provides access to huge collections of distinct, mutually informative descriptions of shared underlying music entities—composers, musical works, performers, and performance recordings. Each provides a degree of machine-readable access to metadata descriptions through an API. However, these metadata are exposed at different levels of granularity and with different capacities for discovering data entities corresponding to specific persons or works—from fairly shallow listings of person names and work titles (in the case of IMSLP) to advanced searches over multiple facets and incorporating external identifiers (in the case of MusicBrainz).

These differences limit the degree to which the different ‘puzzle pieces’ of music information that together ought to contribute to a richer, holistic perspective are *findable*: While each service exposes metadata (*F2*) that explicitly refer (*F3*) to the unique, persistent identifier of the data they describe (*F1*), the degree to which these descriptions are accessible to (and indexed for) external query processes (*F4*) is variable. Though textual labels (person names, work titles, etc.) provide alignment cues in the absence of matching identifiers, they suffer from the ambiguities inherent in such metadata – e.g., names shared among multiple individuals,

works referred to by variant titles, variant spellings, or different languages (Weigl, Lewis, Crawford, Knopke, & Page, 2019).

The data models used to provide metadata descriptions may be customized creations supporting their particular home repository (as in IMSLP), rather than adhering to international standards (as in Muziekweb) (**I2**). Qualified, semantic references between metadata, explicitly exposing the meaning of relationships by specifying, e.g., that the link between a composer and a composition differs from that between a music performer and a composition, may be limited (**I1**, **I3**). This poses further limitations to interoperability between repositories, resulting in a situation where collections such as the ones mentioned here (and there are many, many more!) tend to be used in isolation, their users potentially unaware of further pertinent information available elsewhere.

## 2. CHALLENGES OF IDENTIFICATION

Alongside differences in the way repositories describe their data and expose those descriptions, there is the problem of identifying entities across datasets. Information systems maintaining structured data storage must provide identifiers in order for these data records to be indexed for query and retrieval (**F4**), but unless these identifiers are globally unique and persistent (**F1**), it is not possible to coherently address such data from external sources. Further, unless the descriptive information contained in the identified data records can be exposed (**A1**; taking into account any applicable authentication constraints – **A1.2**), interconnectivity offering holistic, unified views of the data rather than variegated perspectives corresponding to each source repository cannot easily be achieved.

The hypertext transfer protocol (HTTP) alongside its use of uniform resource identifiers (URIs; also extended as internationalized resource identifiers—IRIs) presents an elegant solution to providing and working with globally unique, persistent, resolvable identifiers for Web resources (**A1**, **A1.1**), and also provides the underlying structure of the World Wide Web necessarily used by Web repositories. URIs further provide a means of referring to external entities and thus of providing linkages across repositories.

It is insufficient to merely expose entity identifiers in the hope that these might simply be matched across collections in order to achieve their interconnection, since each collection tends to describe different information facets—the score image on IMSLP is neither the same as the CD described by the Muziekweb catalog, nor the composer identified in MusicBrainz, even when they are strongly related to one another. In order to provide qualified correspondences beyond identity relationships, metadata descriptions must also be exposed using unique, persistent, and dereferenceable identifiers alongside the data they describe (**A1**, **RI**). This is the basis of the Linked Data approach to structured data publication using the Resource Description Framework[10] (RDF). Exposing datasets as Linked Data forms part of best practice around the ‘FAIRification’ process[11] (Jacobsen et al., 2020) of making datasets *Findable*, *Accessible*, *Interoperable* and *Reusable*.

An example of the utility of providing semantic alignment and interlinked, unified access to music library metadata in an early music context is provided by Weigl, Lewis et al. (2019). There, distinct datasets exposing scanned images of music score, machine-readable score encodings, library catalog metadata, and radio broadcast metadata are published as Linked Data according to domain-relevant community standards (**RI.3**), and are interconnected (**I3**), producing a Web of information that allows users to explore episodes of a radio program while inspecting high-resolution score images of presented works, browsing other works by featured composers and their contemporaries, and in turn discovering further radio episodes featuring those works. The resulting nexus of information is published as Linked Data (**F1**, **F3**, **A1.1**), queryable (**F4**) using SPARQL[12], the Web-standard Linked Data query language, and available for public analysis, reuse, and re-interpretation. The process of aligning the multiple datasets involved is complex, requiring both automated algorithms to identify alignment candidates from the vast multitude of possible cross-dataset combinations and drawing on human expertise to confirm or reject these candidates. The inclusion of scholarly insight through provenanced user decisions (**RI.2**, **RI.3**) allow hurdles of historical uncertainty, ambiguity, and scholarly dispute to be addressed by generating inter-dataset relationships modelling not the fact of an alignment, but rather the (citable, contestable) human assertion of one.

### 3. CHALLENGES OF REPRESENTATION

Music information is multimodal by nature; printed music notation and recorded music performances may both offer expressions (through editorial or performative interpretation) of the same musical work, though they are perceived through different senses. Cross-modal analyses are crucial to musicological scholarship (e.g., in the exploration of the range of performative interpretations of a particular piece) and to performance science (e.g., conversely, in the investigation of musical contexts specified in notation that contribute to empirically observed performative effects). They also form a core part of musicianship, notated score acting as an interactive means of reflection upon, and communication about, performative choices (Winget, 2008).

While humans are able to follow (or perform) a score while listening to the music, given that they are suitably enculturated in a particular music tradition and literate in its notation, this remains a problem for machine agents. The problem is particularly tricky where music semantics are not explicitly exposed in the digital representation of a music information resource (**F2**, **RI**); a scanned image of a printed or hand-written piece of notation is, to the computer, merely an otherwise meaningless collection of pixels, an audio recording merely an undifferentiated acoustic signal, without the imposition of algorithmic approximations of meaning derived via optical music recognition and MIR processes.

Music notation represented through explicit machine-readable semantic encoding circumvents this problem, allowing computer processes to work with notation as meaningful content, thus supporting automated processing at a music-theoretical level of analysis, rather than at the level of visual or auditory signals. The Music Encoding Initiative's (MEI) schema for encoding music within a structured, machine-interpretable extensible markup language (XML) representation offers a suitable format (Crawford & Lewis, 2016). Music encodings adhering to the MEI schema comprehensively capture musical meaning within a finely granular, hierarchical XML structure in which every element may be assigned a unique identifier. Paired with the Verovio[13] engraver, which reflects the hierarchy and identifiers of source MEI documents into rendered Scalable Vector Graphic (SVG) representations, this supports the creation of richly interactive Web applications around digital score encodings (Pugin, 2018).

MEI's hierarchical structure of musical meaning, allowing each element to be addressed using a distinct URI (**F1**, **A1**, **A2**), enables the creation of Linked Data to target and richly describe music at varying levels of granularity, from a musical work or movement, to an individual phrase or measure. When MEI encodings are matched to performance recordings, through manual or automated alignment processes, they can further serve as semantic indices into the (otherwise semantically undifferentiated) recorded signal (Lewis, Weigl, & Page, 2019; **F3**, **F4**), providing a basis for FAIR multimedia publishing and communication of musicological materials (Lewis, Weigl, Bullivant, & Page, 2018; **I2**), and for the formulation of semantically enriched digital music objects (Sandler et al., 2019; **II**, **RI**).

### 4. CHALLENGES OF CONTRIBUTION

A key part of the vision of the Web, since its inception, was to enable users around the world to contribute, as well as consume, information (Berners-Lee & Fischetti, 1999). Applying such ideas to the music domain, there is great potential value in encouraging the consumers of music information—scholars, performers, and music enthusiasts—to contribute their expertise, artistic output, and experiential insight to the public sphere. By fostering content creation alongside consumption, we help advance the perpetuation of the classical music tradition, support individual knowledge generation and exchange between users and music information processes, and establish the necessary conditions for projects motivating users to contribute their insight to research as 'citizen scientists' (Jennett & Cox, 2017).

Specialized applications targeting different usage scenarios may contribute to and benefit from shared collections of music information by ensuring that user contributions are ingested and stored in FAIR-compliant fashion (**I2**), with data and metadata organized in structured, machine-readable form (**II**). Thus, for instance, an application catering to music performers by providing access to public-domain music scores, and real-time music information services such as rehearsal-to-score alignment and performance analysis, generates detailed metadata (e.g., regarding onset timings, dynamics, deviations of intonation, or note insertions or omissions) supporting users to track their rehearsal progress over time (Weigl & Goebel, 2020; **RI**).

From a citizen science perspective, such metadata are also promising sources for analysis by performance scientists and music scholars—by collating data derived from different users' rehearsal

renditions (pending their permission – *RI.1*), new kinds of research questions on rehearsal practice based on potentially massive corpora of renditions of the same musical pieces may be investigated.

A data infrastructure supporting contribution can be used to facilitate distributed intelligence, where users' cognitive abilities are harnessed to answer research questions (Haklay, 2012). Depending on the context of research, users may be asked to take concrete training (related to specific actions to be done) and then collect data or carry out simple data analysis and interpretation. This enables large amounts of research data to be created, annotated, analyzed and validated where automated processes are insufficient and human insight is required (*I3, RI, RI.2*). Users, in turn, acquire knowledge and fulfil information needs by participating in such tasks. In order to incentivize users' contributions, it is important to provide feedback relating the value of their contribution and progress to a task, and to transparently reflect up-to-date project results (Bonney et al., 2009). Likewise, users may be engaged to participate when their contributions are acknowledged within a wider user community (Jennett & Cox, 2017; Reeves, Tinati, Zerr, van Kleek, & Simperl, 2017; *RI.1, RI.2*).

In all cases, personally identifiable information contributed by users must be treated with great care for ethical reasons relating to control over personal data and intellectual property, data quality and integrity (Resnik, Elliott, & Miller, 2015), as well as for legal reasons stipulated by the GDPR. Data generated in each context, be it content created for publication by scholars or performers, knowledge exchange between individuals, or participation in wider citizen science projects, necessarily require the collection of personally identifying information, which must be handled in compliance with data protection legislation. To ensure data quality and integrity in citizen science projects, contributors must be trained and monitored to assess and improve their task competence. Collected data and project outcomes should be transparently accessible to participating users and the wider public as far as possible, but sensitive data required for research purposes, such as an individual's gender, birth date, or spoken languages, must be protected (*AI.2, RI.1, RI.2*).

Given these considerations, systems aiming to support organized enrichment of public-domain music repositories while complying with the FAIR principles of accessibility (with associated concerns of access control – *AI.2*) and reusability (with constraints around intellectual property rights – *RI.1, RI.2*) must implement authentication and authorization processes that provide their users with meaningful safeguards of ownership and control over their contributed data.

## 5. CHALLENGES OF RELIABILITY

Automated music processing can be of great value to empirical musicological analysis, providing researchers with precision of measurement and analyses of massive corpora at a scale beyond that achievable by a human listener. Ideally, automatic processing saves time by performing information retrieval (e.g., retrieving key changes from symbolic scores); provides information that a musicologist cannot otherwise extract with satisfactory accuracy (e.g., in determining micro-timings, tempo curves and pitch deviations); and aggregates massive amounts of information extracted from large-scale collections, potentially revealing underlying trends and relationships that may have otherwise eluded research attention.

Practical reality does not currently meet these idealized expectations. MIR technologies can be exploited to automatically derive feature descriptors from music data, approximating meaningful descriptions of melody, tonality, rhythm, structure, emotion or style (*RI*). Such descriptors can be determined at a large scale and complement manual expert analyses. Initiatives such as the MusicBrainz-associated AcousticBrainz project (Porter, Bogdanov, Kaye, Tsukanov, & Serra, 2015), and the Million Song Dataset (Bertin-Mahieux, Ellis, Whitman, & Lamere, 2011) provide audio-derived feature data from massive collections of music content. However, there is still a semantic gap between the provided features (ultimately derived from low-level properties of the acoustic signal) and the higher-level characteristics of relevance in music scholarship (*RI.3*). Further, such datasets are static, describing the outcomes of specific versions of specific algorithms, and may be difficult to reproduce or re-contextualize, e.g., when a new version of an algorithm, or an entirely new MIR technology appears, without the capture of appropriate provenance data (*RI.2*).

## 6. TOWARDS RICHER ONLINE MUSIC PUBLIC-DOMAIN ARCHIVES

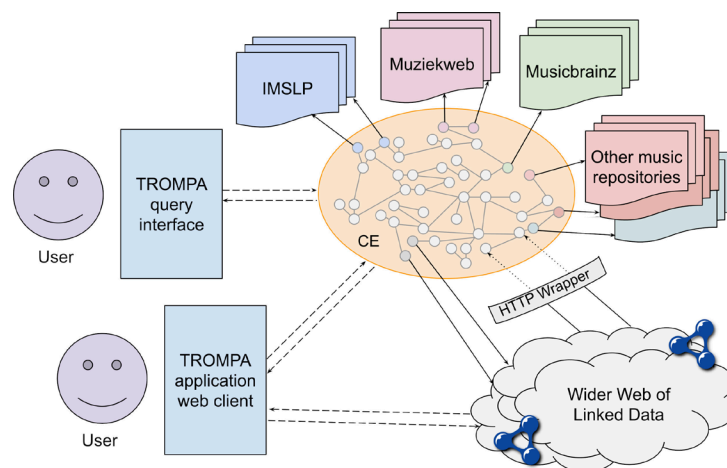
The preceding sections have detailed the challenges faced in applying FAIR principles in the interconnection and enrichment of public-domain music repositories on the Web. We now survey the solutions to these challenges pursued within the context of the TROMPA project. These include: a *data infrastructure* accommodating entities identified within different repositories and described using different data models within a unifying knowledge graph (addressing challenges of *description* and *identification*), *multimodal integration* across different representations of music information through a focus on semantic music encodings and automated alignment techniques (addressing challenges of *representation*), a *federated contribution model* in which user-control over data is retained as far as possible until the moment of explicit publication (addressing challenges of *contribution*), and a focus on the capture of *provenance information* for all machine- and user-initiated data generation processes (addressing challenges of *reliability*).

### 6.1 Data infrastructure

TROMPA’s motivation lies in the interconnection rather than the integration and ingestion of information in public-domain music repositories. It would be costly and counterproductive to attempt to supplant established repositories by copying entity descriptions and media representations into a centralized database under a unified data schema. Rather, we describe the contents of such repositories by reference, using URIs to address, interlink, and contribute layers of enriched descriptors and content to resources hosted *in situ* at their native (TROMPA-external) Web locations (**F1, A1, A2, I3**).

This is achieved by applying a collection of widely used, structured, machine-readable vocabularies (**F2, F3, I1, I2, R1, R1.3**). Schema.org[14], a formalized vocabulary for describing Web resources, provides a core data model for this purpose of virtual data integration across music repositories. We augment this with other well-known, standardized vocabularies. Bibliographic relationships are encoded using the Dublin Core Metadata Initiative’s vocabulary[15]. The Simple Knowledge Organisation System’s[16] (SKOS) vocabulary for mapping relations provides the “glue” interconnecting entities across repositories, as well as connecting multiple metadata records variably describing the same entity within a given repository, as for internationalization. The Web Annotation vocabulary[17] and PROV ontology[18] are used to capture and track the provenance of contributions (**R1.2**) to enrich these resources by TROMPA users and by automated MIR processes. Further established vocabularies are adapted for specialized contexts (**R1.3**), such as for the alignment of musical scores and performance recordings discussed in Section 6.2.

Graph databases are ideally suited to support this flexible, mutably specified interconnection of Web-based resources. TROMPA has opted to adopt a Neo4j property graph database for this purpose (Figure 1).



**Fig. 1.** TROMPA’s Contributor Environment incorporates a neo4j graph database describing music resources housed in external repositories. The HTTP wrapper assigns a URI to each node, exposing it as Linked Data (JSON-LD) when the URI is dereferenced. Dashed arrows: interaction; solid arrows: URI reference; dotted arrows: neo4j to Linked Data translation provided by the HTTP wrapper.



This database, exposed for query via a GraphQL endpoint (*F4*), forms the core of the TROMPA Contributor Environment (CE), a data infrastructure that also comprises a number of component APIs for multimodal query, display, and annotation of music resources, and automated assessment of scores and performances. Querying by SPARQL is not supported; while a database exposing a SPARQL endpoint would allow maximally flexible semantic queries over the CE graph, such endpoints are prone to performance issues at scale (Fields et al., 2015) and do not trivially support the automated processing of newly arriving data driving TROMPA’s enrichment processes (Section 6.4). However, each node in the graph can be accessed via a persistent URI (*F1*) through an HTTP wrapper interface (*A1*), providing a representation of the identified entity and its associated properties and values by reference to their persistent URIs (*F2*, *F3*) using the widely-used JavaScript Object Notation for Linked Data format (JSON-LD[19]). This allows the CE graph to connect with the wider Web of Linked Open Data.

### 6.2 Multimodal integration of music data

Musical score encodings play a key role in several of TROMPA’s specialized user-facing Web applications, as objects of primary interest, but also as musically meaningful scaffolds for the placement of scholarly and performative annotations, and as indices into performance recording timelines.

To accommodate the interconnection of multimodal music information resources by reference to score encodings (*I3*), TROMPA applies a specialized alignment data model implementing the Music Encoding and Linked Data framework (MELD) (Weigl & Page, 2017), combining several standardized, FAIR-compliant vocabularies (*I2*, *RI.3*) to express alignments between performance recordings (Timeline and Event Ontologies; Raimond, Abdallah, Sandler, & Giasson, 2007) and musical score (Segment Ontology; Fields, Page, De Roure, & Crawford, 2011; with anchors in MEI), while interlinking with bibliographic information (Music Ontology; Raimond et al., 2007) and performance feature data (Audio Feature Ontology; Allik, Fazekas, & Sandler, 2016) (see Figure 2).

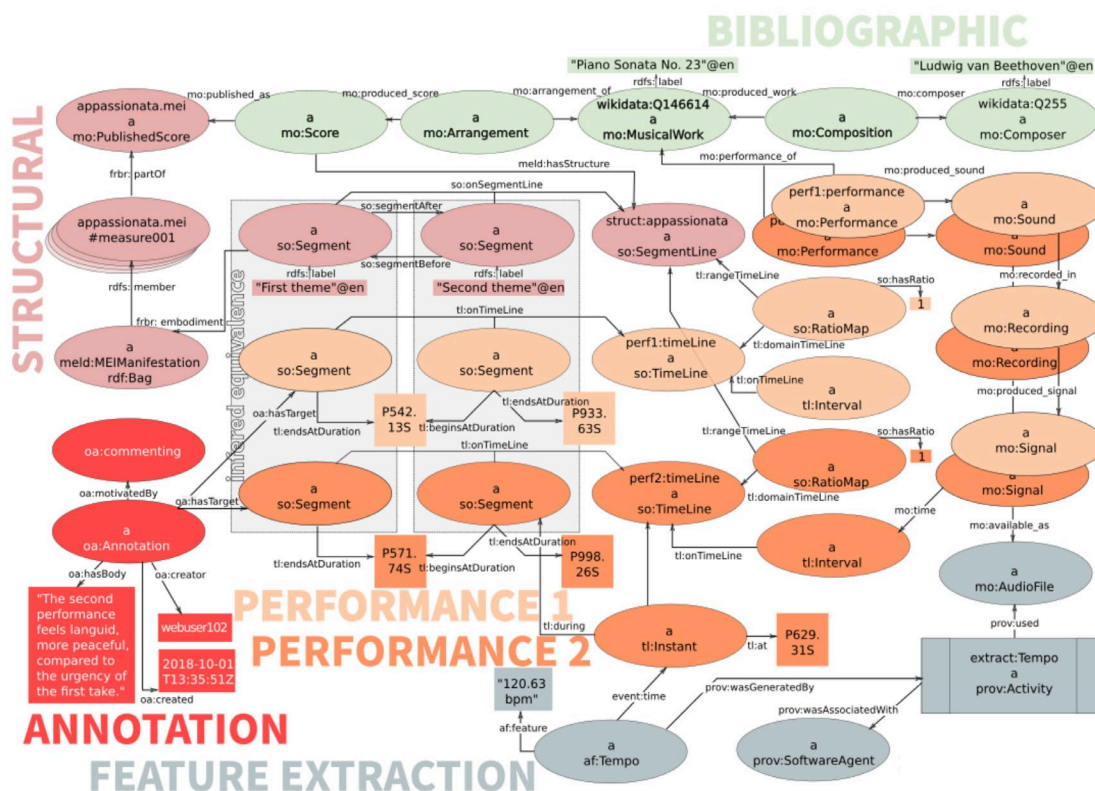
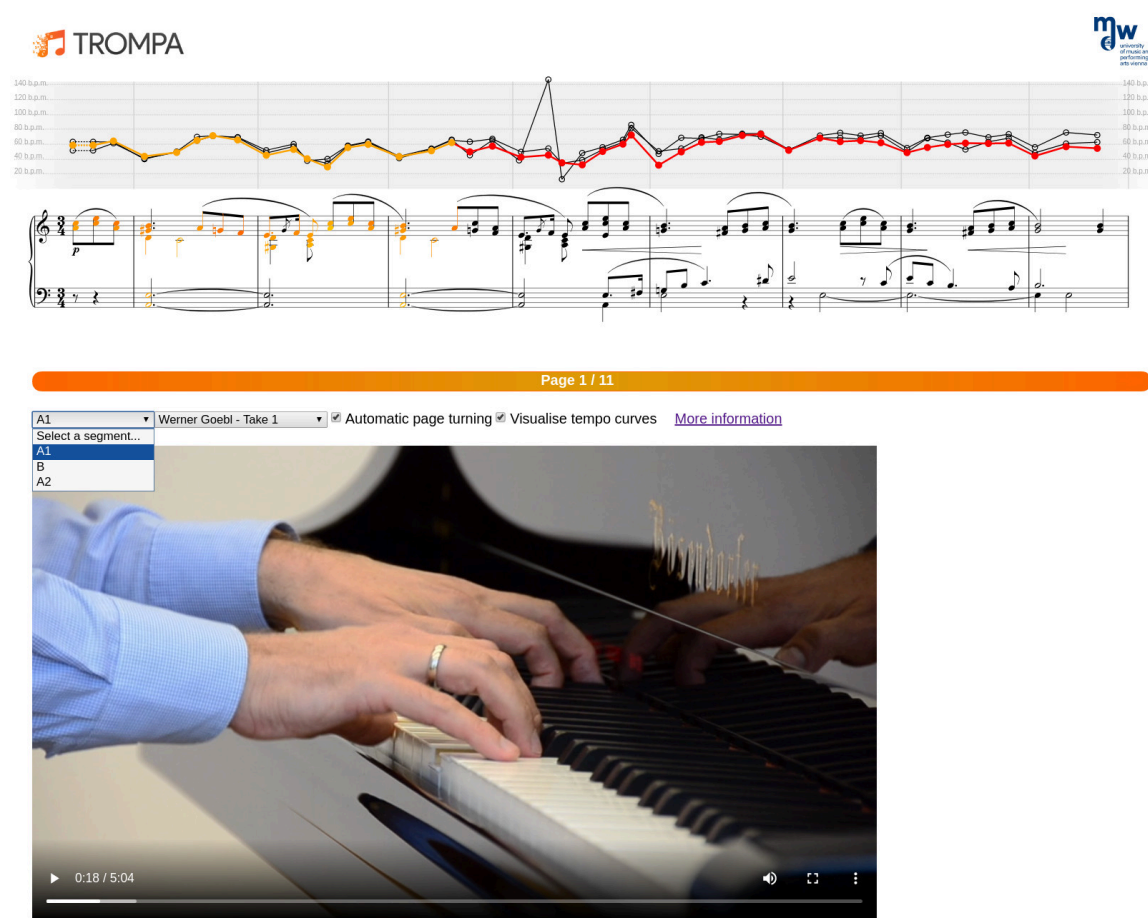


Fig. 2. TROMPA’s alignment data model, implementing the Music Encoding and Linked Data framework (MELD; Weigl & Page, 2017).

The Linked Data representation affords the creation of Web Annotations addressing (fragments of) individual media representations, or indeed their combination in the form of structural segments. Individual entities described using this framework may be published within the TROMPA CE (and exposed by URI through the HTTP wrapper interface) or hosted elsewhere on the Web.

Alignments can be described at different layers of abstraction, e.g., sections, measures, or individual notes. The section-level alignment illustrated in Figure 2 is sufficient for use cases associating performance audio, score information, and aggregated feature data (e.g., average tempo) for salient large-scale musical sections, e.g., to allow users to jump to a particular theme across multimodal representations.

Finer-grained alignments can be expressed by specifying additional segment lines (ordered structural aggregations specified by the Segment ontology), which can be interconnected with segment lines on other hierarchical levels in order to capture highly granular information (e.g., for the analysis of tempo curves or onset timings of individual notes across many recorded renditions). The use of such multilevel alignments is illustrated in the Companion for Long-term Analyses of Rehearsal Attempts (CLARA), TROMPA's web-application targeted at piano players, depicted in Figure 3.



**Fig. 3.** The Companion for Long-term Analyses of Rehearsal Attempts (CLARA[20]; Weigl & Goebel, 2020), one of several specialized TROMPA web applications (available online at <https://trompa.mdw.ac.at>), demonstrates alignment between rendered MEI score encodings, multiple performance timelines, and audiovisual recordings on two levels of granularity: note-level (to generate tempo curves of different performances) and section-level (to generate a navigation drop-down menu allowing users to jump score page and audiovisual playback, here to section A1, B, or A2 of the piece).

The automated determination of alignments between scores and performance recordings is a task subject to continuous refinement in the field of MIR. In the context of TROMPA, we apply tools for real-time and offline alignment employing Hidden Markov Models (Cancino-Chacón et al., 2017; Nakamura,

Ono, Saito, & Sagayama, 2014), extending these to handle MEI encodings through an intermediary step in which the encoded scores are synthesized to MIDI using Verovio. We expose the outputs produced by these tools as Linked Data (*AI*, *AI.2*) conforming to our specialized alignment data model through a custom translation process[21].

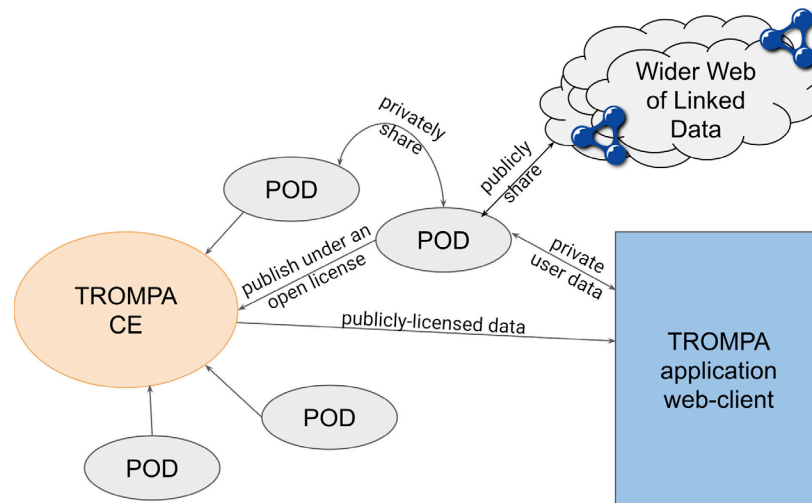
While MEI encodings are highly suitable to the FAIR publication of music information for the reasons discussed above, their availability is limited. Large repositories of public-domain music scores (such as IMSLP, see Section 1) primarily focus on providing digital images (scans) of printed music notation. Commercial composition software such as Sibelius and Finale export MusicXML, an encoding format with an imperative focus on the notation typesetting process, rather than on the declarative description of musical meaning.

To overcome this scarcity of available MEI encodings, TROMPA is pursuing activities to facilitate their creation at scale, employing MIR technologies around optical music recognition, supported and validated using human insight garnered through collaborative crowd-sourcing approaches. The resultant encodings are published under open license (*RI.1*) to repositories hosted within a dedicated GitHub organization[22].

### 6.3 Federated contribution model

Alongside the interconnection and FAIR provision of publicly-licensed music information obtained from established Web repositories, TROMPA's primary motivation lies in the enrichment of such information through the application of MIR technologies and through the contributions of human music scholars, performers, and enthusiasts. Automated enrichment activities are centrally coordinated using the CE alongside the TROMPA Processing Library (see Section 6.4). The resultant data and associated provenance metadata (*RI.2*) are publicly licensed (*RI.1*), becoming available for future reuse and re-interpretation both within TROMPA and in project-external contexts.

While such open publication 'by fiat' (or rather, by the will of the TROMPA consortium) is not particularly problematic where data is generated by automated processing of public-domain information, the situation is complicated considerably in the case of user contributions. Humans reporting on their subjective experiences or providing expert insights or artistic contributions may be understandably concerned about safeguarding their data, and indeed their rights to such safeguards are guaranteed by the EU GDPR and similar legislation. To accommodate these rights, TROMPA's data infrastructure employs a secondary, decentralized layer of so-called personal online datastores (Solid[23] PODs; Figure 4).



**Fig. 4.** TROMPA's federated contribution model allows users to retain data ownership and access control, and to explicitly publish contributions to TROMPA under an open license.

These are user-controlled storage spaces on the Web that allow users and applications with user permission to store data within Web-hosted containers, akin to folders in a file system. Containers and data resources are associated with user-configurable access control lists that enable users to retain control and ownership of their contributions. Solid PODs also act as decentralized identity providers, allowing users to authenticate with TROMPA's applications (Mansour et al., 2016; *AI.2*). Contributions generated by a user's interactions with these applications are stored in the user's POD as Linked Data, referenced by a URI, which can be requested through an HTTP interface (*F1, AI, AI.1*). Users selectively share or retain private access to their generated data items or open up chosen contributions to the public (*AI.2*). PODs may be hosted with any Solid POD provider on the Web, including options for self-hosting by users with the required technological expertise; as such, user-generated contributions hosted in this way are not tied into the core TROMPA infrastructure, and are available for reuse in other contexts and within other applications.

Users may choose to publish their contributions with TROMPA under an open license (Weigl et al., 2020; *RI.1*), at which point the relevant data is ingested into the CE's graph, thus making it discoverable by other TROMPA users and ensuring its persistence within the project even if the copy within the user's POD (or indeed, the POD itself) becomes unavailable (*A2*). Through this mechanism, users are offered fine-grained access control over the information resources they generate through interaction with TROMPA applications, and retain ownership through to the explicit, user-directed act of publication into the public domain. Provenance metadata is captured during the act of publication (*RI.2*), and as the CE exposes its graph as Linked Data, attribution is easily provided using URI references to information stored in the contributing user's POD (*F3*).

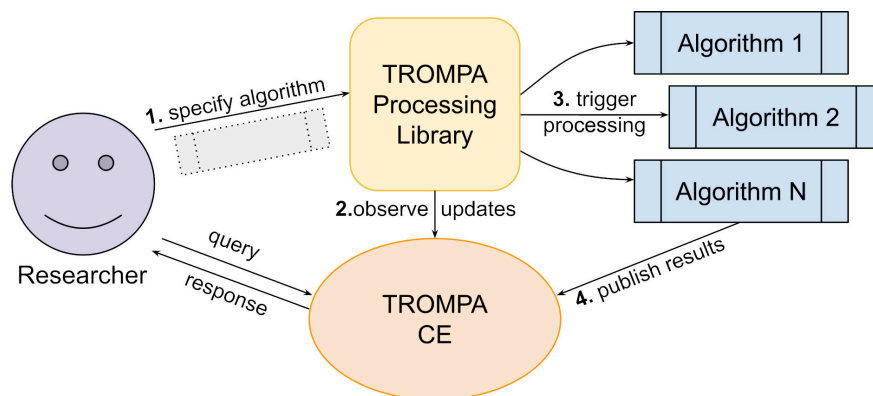
Where users generate data as task responses through participation in a citizen science or crowdsourcing context rather than as self-initiated creative acts, different concerns come into play: Attribution is no longer relevant where such data is collected anonymously, but sensitive data may need to be retained (and thus safeguarded) for analysis by the researchers directing the citizen science project. Listening and perception studies conducted through this platform must capture sensitive information if demographic analyses, e.g., according to gender, cultural background, or age are to be performed. Likewise, anonymized contributions retaining some sensitive information are highly valued for designing customized incentivization mechanisms such as musical recommendations, task performance feedback or community and personal analytics. Sensitive data obtained within this context are managed only by authorized researchers within the TROMPA consortium. Datasets are comprehensively anonymized before publication, and then shared with the wider research community through the Zenodo[24] open science platform. In any such case, GDPR principles are followed, TROMPA's ethical procedure for research has been approved by the Committee for Ethical Review of consortium members, and participants are always informed about the management of their personal data, about the scope and goals of the citizen science project in which they chose to participate, and about the types of result that will be shared with the research community.

## 6.4 Data generation and provenance capture

The outcomes of data-generating processes within the TROMPA project are associated with provenance descriptors (*RI.2*) and explicit licensing information (*RI.1*) in order to provide adequate means to reproduce, re-interpret and re-contextualize data in future use, both within the project and in external contexts. For user contributions, this is usually done at the point of publication, referencing profile information within a user's POD where attribution is required (*F3*).

TROMPA's automated processing employs established state-of-the-art algorithms, as well as new MIR technologies developed in the course of the project. These algorithms are packaged within the TROMPA Processing Library (TPL), a conceptual organization of software algorithms (see Figure 5) that integrates with the CE through an HTTP interface (*AI.1*), orchestrating automated processing of information resources referenced by the knowledge graph.

The TPL accommodates the specification of algorithms alongside provenance information relating to their processing context, including software versions and parameter values. It also provides a subscription mechanism by which algorithms may be automatically triggered to process newly created nodes matching specified type constraints as they are added to the CE's knowledge graph. References to processing outcomes (generally MIR feature data) are themselves ingested into the knowledge graph (*F1, F3, F4*), where they may trigger further activities orchestrated by the TPL, resulting in processing chains. These mechanisms allow researchers to modularly specify new algorithms or new software versions as they become available, and to run them on demand in response to the arrival of specific types of data.



**Fig. 5.** TROMPA Processing Library (TPL) workflow. (1.) The researcher specifies their algorithm for use with the TPL. (2.) The TPL continuously monitors the CE’s graph. When a new node is ingested with a type matching a TPL algorithm’s specification, the algorithm is triggered on the newly entered data (3.), before publishing the results to the CE (4.), where they become available for query and may potentially trigger further processing orchestrated by the TPL.

In order to be approved for inclusion within the TPL, software components must fulfil a number of requirements to ensure the suitability of their output within an open science context. To support reproducibility, software components as well as their dependencies should be openly licensed, with source code deposited online (e.g., on GitHub). The algorithms implemented by software components should be well-documented and ideally reference scientific publications that adequately describe the algorithm. The software components themselves should also be documented, providing clear instructions on how the software may be set up and run outside a TROMPA context, to ensure capacity for external re-use and verification of results. Finally, all software should ideally be packaged alongside their dependencies within a modular executable processing context (e.g., via a Docker container; as per Cito, Ferme, and Gall, 2016) to simplify reproduction and verification of results while minimizing constraints of technical expertise.

## 7. CONCLUSION

We have presented TROMPA, a project aiming to facilitate engagement with public-domain music repositories available on the Web, interlinking and exposing music resources to allow them to serve many different potential use cases. The project embraces user contributions generated in these use cases, which are captured and reintegrated in the form of further interlinked connections.

Our vision is to democratize the usage of musical resources: Amateurs and professionals, players and scholars may all wish to engage with the same music object, their divergent perspectives and different approaches to description and use contributing to a more comprehensive conception of the object itself. We acknowledge and respect this diversity of contributions by enabling different interlinked musical representations to co-exist, by establishing a federated contribution model safeguarding explicit user control of their data and by focusing on the capture of provenance information in all data-generating activities.

We have illustrated challenges and presented solutions toward implementing the data infrastructure and music information services supporting this vision in compliance with FAIR principles and Web standards to ensure that such broad reuse will indeed be possible. With competing project motivations around attribution and protection of intellectual contributions on the one hand and a mission to capture and generate massive amounts of openly licensed data on the other, concerns around data and user management pose non-trivial but very timely and interesting research problems. We strongly believe that TROMPA as a publicly funded European project can pioneer means of addressing challenges in this area in a GDPR-compliant fashion with explicit attention to the public cause.

Technological considerations behind TROMPA’s data infrastructure and approach to music information may initially seem far removed from more common musicological interests, which may rather be geared towards specific musical repertoires. We do not seek to establish new corpora of our own; rather, we want to help make existing corpora (which have existing champions and maintainers) FAIRer, richer, and

more interconnected. In this way, musicological insights into dedicated repertoires may become more broadly accessible.

At a more abstract level, musicological scholarship typically employs considerable representational and interpretational sophistication, often demanding for customizable, personalizable and contestable perspectives to be added to multifaceted music objects at a much finer-grained level than typically would be seen in data on the Web. TROMPA's infrastructure and representational models are explicitly informed by these scholarly requirements. We hope that TROMPA's perspective will benefit and strengthen open data and its infrastructures at large.

## ACKNOWLEDGEMENTS

The TROMPA Project is funded under the European Union's Horizon 2020 research and innovation programme *H2020-EU.3.6.3.1. - Study European heritage, memory, identity, integration and cultural interaction and translation, including its representations in cultural and scientific collections, archives and museums, to better inform and understand the present by richer interpretations of the past* under grant agreement No 770376. We gratefully acknowledge the contributions of our colleagues in the TROMPA consortium. This article was copyedited by Christine Ahrends and layout edited by Diana Kayser.

## NOTES

[1] Correspondence can be addressed to: Dr. David M. Weigl, Dept. of Music Acoustics – Wiener Klangstil, University of Music and Performing Arts Vienna, Anton-von-Webern-Platz 1/II, 1030 Vienna, Austria. E-mail: weigl@mdw.ac.at

[2] TROMPA Project website available at <https://trompamusic.eu>

[3] International Music Score Library Project / Petrucci Music Library available at <https://imslp.org>

[4] IMSLP API documentation available at <https://imslp.org/wiki/IMSLP:API>

[5] Muziekweb website available at <https://www.muziekweb.eu>; Muziekweb is an institutional member of the TROMPA consortium.

[6] OCLC Bibliographic Formats and Standards description available at <https://www.oclc.org/bibformats/en.html>

[7] Muziekweb API documentation available at <https://www.muziekweb.nl/Muziekweb/Webservice/WebserviceAPI.php>

[8] An archived description of the BBC's use of MusicBrainz is available at <http://www.bbc.co.uk/music/brainz/>

[9] Refer to descriptions of the FAIR principles at <https://www.go-fair.org/fair-principles/> and to Wilkinson et al. (2016)

[10] RDF breaks down data relationships of arbitrary complexity into minimal meaningful components, in the form of (subject, predicate, object) triples. Each component of a triple may take the form of a URI, supporting the chaining of triples by placing the URI that is the object of one triple into the subject position of another triple; and through this mechanism, enabling the creation of knowledge graphs. Refer to the RDF primer at <http://www.w3.org/TR/rdf11-primer>

[11] Best practices for the FAIRification process are described at <https://www.go-fair.org/fair-principles/fairification-process/>

[12] SPARQL is a standard language for querying graphs of Linked Data; refer to <https://www.w3.org/TR/sparql11-query/>

[13] Verovio is a command-line tool, Python module, and Javascript library for turning MEI encodings into beautifully rendered music scores. Refer to the Verovio website at <https://www.verovio.org>

[14] Schema.org is a standard vocabulary for expressing structured data about Web resources. Refer to <https://schema.org/>

[15] Refer to DCMI documentation at <https://dublincore.org/specifications/dublin-core/dcmi-terms/>

[16] The Simple Knowledge Organization System (SKOS) is a common data model for linking knowledge organization systems via the Semantic Web. Refer to SKOS primer at <https://www.w3.org/2009/08/skos-reference/skos.html#SKOS-PRIMER>

[17] Web Annotations are a standard vocabulary and data model for expressing annotations about Web resources. Refer to <https://www.w3.org/TR/annotation-vocab/>

[18] The PROV Ontology is a standard specification of classes, properties, and restrictions used to represent and exchange provenance information. Refer to the overview of PROV at <https://www.w3.org/TR/prov-overview/>

[19] JSON-LD is a serialization format for Linked Data (RDF). It is especially convenient for software developers, as it is a compatible specialization of the standard JSON format widely used to capture and exchange data structures in automated processing. Refer to <https://www.w3.org/TR/json-ld/>

[20] The Companion for Long-term Analysis of Rehearsal Attempts (CLARA), a TROMPA web-application available at <https://trompa.mdw.ac.at>, providing instrumental players with insight into their rehearsal practice. Source code available under open license from <https://github.com/trompamusic/clara>. CLARA is named in honor of Clara Schumann, on whose 200<sup>th</sup> centenary the software was initially released.

[21] MEI-anchoring and RDF conversion scripts available under open license from <https://github.com/trompamusic/trompa-align>

[22] MEI encodings generated by activities within the TROMPA project are available under open license from <https://github.com/trompamusic-encodings>

[23] Solid is a web decentralization project with a strong focus on user-retained control over data. Refer to project website at <https://solidproject.org>

[24] The TROMPA Zenodo community can be found at <https://zenodo.org/communities/trompa>

## REFERENCES

Allik, A., Fazekas, G., & Sandler, M. B. (2016). An ontology for audio features. In J. Devaney, M. I. Mandel, D. Turnbull, & G. Tzanetakis (Eds.), *Proceedings of the 17th International Society for Music Information Retrieval Conference*, New York City, NY, USA (pp. 73–79). <http://archives.ismir.net/ismir2016/paper/000077.pdf>

Berners-Lee, T., & Fischetti, M. (1999). *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*, New York, NY: Harper Collins.

Bertin-Mahieux, T., Ellis, D. P. W., Whitman, B., & Lamere, P. (2011). The million song dataset. In A. Klapuri, & C. Leider (Eds.), *Proceedings of the 12th International Society for Music Information Retrieval Conference*, Miami, FL, USA (pp. 591–596). <https://archives.ismir.net/ismir2011/paper/000022.pdf>

Bonney, R., Cooper, C. B., Dickinson, J., Kelling, S., Phillips, T., Rosenberg, K. V., & Shirk, J. (2009). Citizen science: A developing tool for expanding science knowledge and scientific literacy. *BioScience*, 59(11), 977–984. <https://doi.org/10.1525/bio.2009.59.11.9>

Cancino-Chacón, C., Bonev, M., Durand, A., Grachten, M., Arzt, A., Bishop, L., ... Widmer, G. (2017). The ACCompanion v0.1: An expressive accompaniment system. *Late Breaking / Demos session, 18th International Society for Music Information Retrieval Conference*, Suzhou, China (pp. 1–2). [https://ui.adsabs.harvard.edu/link\\_gateway/2017arXiv171102427C/arxiv:1711.02427](https://ui.adsabs.harvard.edu/link_gateway/2017arXiv171102427C/arxiv:1711.02427)

Cito, J., Ferme, V., & Gall, H. C. (2016). Using Docker containers to improve reproducibility in software and web engineering research. In A. Bozzon, P. Cudre-Maroux, & C. Pautasso (Eds.), *Proceedings of the International Conference on Web Engineering*, Lugano, Switzerland (pp. 609–612). Cham, Switzerland: Springer. [https://doi.org/10.1007/978-3-319-38791-8\\_58](https://doi.org/10.1007/978-3-319-38791-8_58)

Crawford, T., & Lewis, R. (2016). Review: Music Encoding Initiative. *Journal of the American Musicological Society* 69(1), 273–285. <https://doi.org/10.1525/jams.2016.69.1.273>

Fields, B., Page, K., De Roure, D., & Crawford, T. (2011). The segment ontology: Bridging music-generic and domain-specific. In *Proceedings of the 2011 IEEE International Conference on Multimedia and Expo*, Barcelona, Spain (pp. 1–6). New York, NY, USA: IEEE. <https://doi.org/10.1109/icme.2011.6012204>

Fields, B., Phippen, P., & Cohen, B. (2015). A case study in pragmatism: Exploring the practical failure modes of Linked Data as applied to classical music catalogues. In B. Fields & K. Page (Eds.), *Proceedings of the 2nd International Workshop on Digital Libraries for Musicology*, Knoxville, TN, USA (pp. 21–24). New York, NY, USA: ACM. <https://doi.org/10.1145/2785527.2785531>

Haklay, M. (2012). Citizen science and volunteered geographic information: Overview and typology of participation. In D. Z. Sui, S. Elwood, & M. F. Goodchild (Eds.), *Crowdsourcing Geographic Knowledge* (pp. 105–122). [https://doi.org/10.1007/978-94-007-4587-2\\_7](https://doi.org/10.1007/978-94-007-4587-2_7)

Hasnain, A., & Rebholz-Schuhman, D. (2018). Assessing FAIR data principles against the 5-Star open data principles. In A. Gangemi, A. L. Gentile, A. G. Nuzzolese, S. Rudolph, M. Maleshkova, H. Paulheim, ... M. Alam (Eds.), *The Semantic Web: ESWC 2018 Satellite Events. ESWC 2018. Lecture Notes in Computer Science* (Vol. 11155) (pp. 469–477). Cham, Switzerland: Springer. [https://doi.org/10.1007/978-3-319-98192-5\\_60](https://doi.org/10.1007/978-3-319-98192-5_60)

Jacobsen, A., Kaliyaperumal, R., da Silva Santos, L. O. B., Mons, B., Schultes, E., Roos, M., & Thompson, M. (2020). A generic workflow for the data FAIRification process. *Data Intelligence* 2(1–2), 56–65. [https://doi.org/10.1162/dint\\_a\\_00028](https://doi.org/10.1162/dint_a_00028)

Jennett, C., & Cox, A. L. (2017). Digital citizen science and the motivations of volunteers. In K. L. Norman, & J. Kirakowski (Eds.), *The Wiley Handbook of Human Computer Interaction* (pp. 831–841). <https://doi.org/10.1002/9781118976005.ch39>

Lewis, D., Weigl, D. M., Bullivant, J., & Page, K. R. (2018). Publishing musicology using multimedia digital libraries: Creating interactive articles through a framework for Linked Data and MEI. In K. Page (Ed.), *Proceedings of the 5th International Conference on Digital Libraries for Musicology*, Paris, France (pp. 21–25). New York, NY, USA: ACM. <https://doi.org/10.1145/3273024.3273038>

Lewis, D., Weigl, D. M., & Page, K. R. (2019). Musicological observations during rehearsal and performance: A Linked Data digital library for annotations. In D. Rizo (Ed.), *Proceedings of the 6th International Conference on Digital Libraries for Musicology*, The Hague, The Netherlands (pp. 1–8). New York, NY, USA: ACM. <https://doi.org/10.1145/3358664.3358669>

Lewis, R. J., Crawford, T., & Lewis, D. (2015). Exploring information retrieval, semantic technologies and workflows for music scholarship: The Transforming Musicology project. *Early Music*, 43(4), 635–647. <https://doi.org/10.1093/em/cav073>



- Liem, C. C. S., Gómez, E., & Schedl, M. (2015). PHENICX: Innovating the classical music experience. In *2015 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)* (pp. 1–4). New York, NY, USA: IEEE. <https://doi.org/10.1109/ICMEW.2015.7169835>
- Liem, C. C. S., Müller, M., Eck, D., Tzanetakis, G., & Hanjalic, A. (2011). The need for music information retrieval with user-centered and multimodal strategies. In C. C. S. Liem, M. Müller, D. Eck, & G. Tzanetakis (Eds.), *Proceedings of the 1st International ACM Workshop on Music Information Retrieval with User-centered and Multimodal Strategies (MIRUM '11)*, Scottsdale, AZ, USA (pp. 1–6). New York, NY, USA: ACM. <https://doi.org/10.1145/2072529.2072531>
- Mansour, E., Samba, A. V., Hawke, S., Zereba, M., Capadisli, S., Ghanem, A., ... Berners-Lee, T.. (2016). A demonstration of the Solid platform for social web applications. In J. Bourdeau, J. A. Hendler, & R. N. Nkambou (Eds.), *Proceedings of the 25th International Conference Companion on World Wide Web*, Montreal, QC, Canada (pp. 223–226). <https://doi.org/10.1145/2872518.2890529>
- Nakamura, E., Ono, N., Saito, Y., & Sagayama, S. (2014). Merged-output Hidden Markov Model for score following of MIDI performances with ornaments, desynchronized voices, repeats and skips. In K. Moschos (Ed.), *Proceedings of the 2014 Joint SMC/ICMC Conference*, Athens, Greece (pp. 1185–1192). <http://smc.afim-asso.org/smc-icmc-2014/images/proceedings/OS19-B08-Merged-OutputHidden.pdf>
- Page, K. R., Bechhofer, S., Fazekas, G., Weigl, D. M., & Wilmering, T. (2017). Realising a layered digital library: Exploration and analysis of the live music archive through Linked Data. In R. H. McDonald, & N. Worby (Eds.), *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, Toronto, ON, Canada (pp. 1–10). New York, NY, USA: IEEE. <https://doi.org/10.1109/JCDL.2017.7991563>
- Porter, A., Bogdanov, D., Kaye, R., Tsukanov, R., & Serra, X. (2015). AcousticBrainz: A community platform for gathering music information obtained from audio. In M. Müller, & F. Wiering (Eds.), *Proceedings of the 16th International Society for Music Information Retrieval Conference*, Málaga, Spain (pp. 786–792). <https://archives.ismir.net/ismir2015/paper/000210.pdf>
- Pugin, L. (2018). Interaction perspectives for music notation applications. In S. Bechhofer, G. Fazekas, & K. Page (Eds.), *Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music*, Monterey, CA, USA (pp. 54–58). New York, NY, USA: ACM. <https://doi.org/10.1145/3243907.3243911>
- Raimond, Y., Abdallah, S. A., Sandler, M. B., & Giasson, F. (2007). The Music Ontology. In S. Dixon, D. Bainbridge, & R. Tybke (Eds.), *Proceedings of the 8th International Society for Music Information Retrieval Conference*, Vienna, Austria (pp. 1–6). [https://ismir2007.ismir.net/proceedings/ISMIR2007\\_p417\\_aimond.pdf](https://ismir2007.ismir.net/proceedings/ISMIR2007_p417_aimond.pdf)
- Reeves, N., Tinati, R., Zerr, S., van Kleek, M. G., & Simperl, E. (2017). From crowd to community: A survey of online community features in citizen science projects. In C. P. Lee, & S. Poltrock (Eds.), *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17*, Portland, OR, USA (pp. 2137–2152). New York, NY, USA: ACM. <https://doi.org/10.1145/2998181.2998302>
- Resnik, D. B., Elliott, K. C., & Miller, A. K. (2015). A framework for addressing ethical issues in citizen science. *Environmental Science & Policy*, *54*, 475–481. <https://doi.org/10.1016/j.envsci.2015.05.008>
- Sandler, M., De Roure, D., Benford, S., & Page, K. R. (2019). Semantic web technology for new experiences throughout the music production-consumption chain. In A. Baratè, L. A. Ludovico, S. Ntalampiras, & G. Presti (Eds.), *Proceedings of the 2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*, Milano, Italy (pp. 49–55). New York, NY, USA: IEEE. <https://doi.org/10.1109/mmrp.2019.00017>
- Swartz, A. (2002). MusicBrainz: A semantic web service. *IEEE Intelligent Systems*, *17*(1), 76–77. <https://doi.org/10.1109/5254.988466>

- van Reisen, M., Stokmans, M., Basajja, M., Ong'ayo, A. O., Kirkpatrick, C., & Mons, B. (2020). Towards the tipping point for FAIR implementation. *Data Intelligence*, 2(1-2), 264–275. [https://doi.org/10.1162/dint\\_a\\_00049](https://doi.org/10.1162/dint_a_00049)
- Voigt, P., & von dem Bussche, A. (2017). *The EU General Data Protection Regulation (GDPR). A practical guide* (1st ed.). Cham, Switzerland: Springer. <https://doi.org/10.1007/978-3-319-57959-7>
- Weigl, D. M., & Page, K. R. (2017). A framework for distributed semantic annotation of musical score: “Take it to the bridge!”. In X. Hu, S. J. Cunningham, D. Turnbull, & Z. Duan (Eds.), *Proceedings of the 18th International Society for Music Information Retrieval Conference*, Suzhou, China (pp. 221–228). <http://archives.ismir.net/ismir2017/paper/000190.pdf>
- Weigl, D. M., Goebel, W., Crawford, T., Gkiokas, A., Gutierrez, N. F., Porter, A., ... van Tilburg, M. (2019). Interweaving and enriching digital music collections for scholarship, performance, and enjoyment. In D. Rizo (Ed.), *Proceedings of the 6th International Conference on Digital Libraries for Musicology*, The Hague, The Netherlands (pp. 84–88). New York, NY, USA: ACM. <https://doi.org/10.1145/3358664.3358666>
- Weigl, D. M., Lewis, D., Crawford, T., Knopke, I., & Page, K. R. (2019). On providing semantic alignment and unified access to music library metadata. *International Journal on Digital Libraries*, 20(1), 25–47. <https://doi.org/10.1007/s00799-017-0223-9>
- Weigl, D. M., & Goebel, W. (2020). Rehearsal encodings with a social life. In E. De Luca, & A. J. Kijas (Eds.), *Music Encoding Conference Proceedings 2020*, Boston, MA, USA (online) (pp. 51–53). <https://doi.org/10.17613/5ae5-8387>
- Weigl, D. M., Goebel, W., Hofmann, A., Crawford, T., Zubani, F., Liem, C. C. S., & Porter, A. (2020). Read/write digital libraries for musicology. In D. Lewis (Ed.), *Proceedings of the 7th International Conference on Digital Libraries for Musicology*, Montreal, QC, Canada (online) (pp. 48–52). New York, NY, USA: ACM. <https://doi.org/10.1145/3424911.3425519>
- Wiggins, G. A. (2016). Computer representation of music in the research environment. In T. Crawford & L. Gibson (Eds.), *Modern Methods for Musicology* (pp. 27–42). Abingdon, UK: Routledge. <https://doi.org/10.4324/9781315595894-12>
- Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., ... Mons, B. (2016). The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3, 160018. <https://doi.org/10.1038/sdata.2016.18>
- Winget, M. A. (2008). Annotations on musical scores by performing musicians: Collaborative models, interactive methods, and music digital library tool development. *Journal of the Association for Information Science and Technology* 59(12), 1878–1897. <https://doi.org/10.1002/asi.20876>