Interweaving and enriching digital music collections for scholarship, performance, and enjoyment

Weigl, David M.; Goebl, Werner; Crawford, Tim; Gkiokas, Aggelos; Gutierrez, Nicolas F.; Porter, Alastair; Santos, Patricia; Karreman, Casper; Vroomen, Ingmar; Liem, Cynthia C.S.

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Interweaving and Enriching Digital Music Collections for Scholarship, Performance, and Enjoyment

David M. Weigl, Werner Goebl
{lastname}@mdw.ac.at
University of Music and Performing Arts Vienna

Tim Crawford
t.crawford@gold.ac.uk
Goldsmiths, University of London

Aggelos Gkiokas, Nicolas F. Gutierrez, Alastair Porter, Patricia Santos
{firstname.lastname}@upf.edu
Universitat Pompeu Fabra

Casper Karreman, Ingmar Vroomen
{ckarreman|ingmar}@muziekweb.nl
Muziekweb

Cynthia C. S. Liem
C.C.S.Liem@tudelft.nl
Delft University of Technology

Álvaro Sarasúa
alvaro.sarasua@voctrolabs.com
Voctro Labs S.L.

Marcel van Tilburg
m.vtilburg@concertgebouworkest.nl
Royal Concertgebouw Orchestra

ABSTRACT

The turn toward the digital has opened up previously difficult to access musical materials to wider musicological scholarship. Digital repositories provide access to publicly licensed score images, score encodings, textual resources, audiovisual recordings, and music metadata. While each repository reveals rich information for scholarly investigation, the unified exploration and analysis of separate digital collections remains a challenge. TROMPA—*Towards Richer Online Music Public-domain Archives*—addresses this through a knowledge graph interweaving composers, performers, and works described in established digital music libraries, facilitating discovery and combined access of complementary materials across collections. TROMPA provides for contribution of expert insights as citable, provenanced annotations, supporting analytical workflows and scholarly communication. Beyond scholars, the project targets four further user types: instrumental players; choir singers; orchestras; and music enthusiasts; with corresponding web applications providing specialised views of the same underlying knowledge graph. Thus, scholars' annotations provide contextual information to other types of users; while performers' rehearsal recordings and performative annotations, conductors' marked up scores, and enthusiasts' social discussions and listening behaviours, become available to scholarly analysis (per user consent). The knowledge graph is exposed as Linked Data, adhering to the FAIR principles of making data Findable, Accessible, Interoperable, and Re-usable, and supporting further interlinking, re-interpretation and re-use beyond the immediate scope of the project.

CCS CONCEPTS

• **Information systems** → **Digital libraries and archives**; **Music retrieval**; *Crowdsourcing*; *Web services*.

KEYWORDS

Public domain, music archives, data infrastructure, linked data

1 MOTIVATION: ENRICHING CLASSICAL MUSIC

Classical music represents treasured cultural heritage and contemporary tradition, perpetuated and re-interpreted through practice, performance, scholarship, and enjoyment. Music libraries and archives assemble, preserve, and organise classical music resources, but underserve our more dynamic interactions with this repertoire. Enriching these interactions is important to engage and diversify the classical music audience, thus sustaining this tradition [18].

TROMPA[1]—*Towards Richer Online Music Public-domain Archives*—an EU-funded three year project currently completing its first year of activity, addresses this challenge by combining music information retrieval (MIR) technologies and crowd-sourcing approaches to publish, interlink, contextualise, and augment public-domain classical music resources. Building on established music repositories, TROMPA provides for discovery, enhancement, and contribution of musical scores, recordings, analyses, and interpretations, applying standard Web- and MIR-technologies to ensure reusable, scalable, and sustainable access to the data produced.

[1]https://trompamusic.eu

2 PRECEDING PROJECTS

TROMPA builds on a number of related research projects on inter-linking and enrichment of music information resources.

Performances as Highly Enriched aNd Interactive Concert eXperiences (PHENICX). PHENICX was an international research project funded under the EU's Seventh Framework Programme. Motivated by notions that technological developments in the current digital age could offer new opportunities to make symphonic classical music more accessible to broader audiences, the project had two main focus areas [13, 16]. Research was performed both into improving audiovisual analysis techniques necessary for enabling multimodal enrichment; and, into finding ways to make such enrichments engaging and useful to the intended broader audiences.

Though producing impactful demonstrations of technologically enriched concert experiences (e.g. [1, 11, 19, 23]), scalability was limited. Professional-level enrichment requires the availability of clean and well-structured input data; in practice, this required considerable and expensive involvement of human experts, even though much of the needed human work would consist of quality checks. TROMPA addresses this limitation through crowd-sourcing components scalably involving human insight in enrichment activities.

Transforming Musicology. This wide-ranging UK-based, AHRC-funded project included a focus on semantic linking of musical resources and workflows, demonstrating how scholars might take fuller advantage of the possibilities for presentation, analysis and discovery inherent in a web of digital resources organised as Linked Data. One complex case study captured bio-physical reactions of 10 audience members in live performances of Wagner's *Ring* cycle, with real-time annotations of staging details, resulting in a multi-modal, time-aligned dataset in which data could be displayed or analysed in synchrony with musical score or recorded audio [21].

Methods for capturing the semantics of scholarly workflows were studied, analysing and comparing steps needed to achieve useful results in music(ologic)al tasks [20]. Other work, in the SLICK-MEM and SLoBR projects immediately preceding Transforming Musicology [6, 26], focussed on the problems of aligning multiple datasets compiled with inconsistent formats or standards. Another study around querying interlinked classical music catalogues [9] exposed limitations of SPARQL, the query language in general use for searching Linked Data. These will be addressed in TROMPA by the use of a graph-based data infrastructure in which searching is performed with the more scalability-oriented GraphQL language.

Fusing Audio and Semantic Technologies (FAST). A UK EPSRC funded project currently concluding 5 years of research on audio processing technologies, studio science, and the Semantic Web, FAST defines Digital Music Objects (DMOs), flexible constructs coupling recorded music essence with rich, semantic, linked metadata [22], and explores their application throughout the music value chain, from production, through distribution, to consumption. DMOs retain provenance traces of their activities throughout this chain, with implications for digital libraries for musicology [7]. The DMO notion is particularly informative in TROMPA's environment featuring the interlinking of music metadata, provenance-tracked contributions by human and machine agents, and re-use and re-interpretation within different usage contexts.

One output of the FAST project finding direct re-application within TROMPA is MELD (*Music Encoding and Linked Data*) [27], a semantic framework and open-source toolkit for the creation of dynamic digital scores incorporating Web Annotations[2]. FAST applied MELD in music performance [14] and (re-)composition [8], and it has been used as a means of publishing musicology in a multimedia digital library context [15] outside of the project. Within TROMPA, MELD forms the basis of the digital score edition component, where it will be used to facilitate the creation and display of scholarly, performative, and discursive score annotations.

3 PUBLIC DOMAIN MUSIC DL RESOURCES

Numerous repositories offer public-domain classical music resources on the Web, with audio (and audio-visual) materials, score encodings, digitised sheet music and other printed material like books and articles, sources with contextual and historical background information and metadata. These sources are published by (music) archives and libraries, broadcasters, universities, research institutions and scholars, but also by commercial companies or individuals.

One of the largest and most notable collections, the International Music Score Library Project (IMSLP), also known as the Petrucci Music Library, [3] contains over 475,000 scores by more than 17,500 composers. All scores included in IMSLP belong to the public domain in either Canada or the US. The IMSLP is an important source for musicians and scholars seeking printed editions of classical music pieces, often offering multiple versions of the same composition. IMSLP also contains Creative Commons-licensed recordings uploaded by users, and links to commercial recordings provided by music labels, which paid subscribers can listen to.

Another important public-domain classical music score repository is the Choral Public Domain Library (CPDL), which holds over 32,000 choral and vocal works by at least 3,200 composers.[4] Both IMSLP and CPDL are important repositories as sources for different technologies and use cases in the TROMPA project.

The main resource for public domain structured (and machine-readable) music metadata is MusicBrainz, an "open music encyclopedia" maintained by a global community of users.[5] Although aimed broadly at music of all genres, MusicBrainz contains an impressive number of classical works, composers and performers. The MusicBrainz data model includes many features that uniquely suit classical music, including distinctly identifying compositions and movements, annotating compositions with catalogue numbers, and relating recordings to people who participated in them—e.g., performing orchestra, any soloists, the conductor—as well as specific information about composers and works performed. Data quality and quantity vary on initial contribution, but community members can correct, adjust or complement the data. MusicBrainz' structured data model and use of unique identifiers have made it an authority for music identification, supported by external links to other music repositories, websites and streaming services.

The biggest non-commercial collections of audio recordings can be found in specialised music archives and libraries. They are often part of national libraries, like the British Library Sound Archive

---

[2]https://www.w3.org/TR/annotation-model
[3]https://imslp.org
[4]https://www.cpdl.org/wiki
[5]https://musicbrainz.org

or the *Deutsches Musikarchiv*.[6] Such collections are generally not publicly accessible outside of their source institution, remaining effectively 'invisible' (no audio playback; no display of artwork or record covers) and not searchable without specialist access.

Muziekweb, based in Rotterdam, does provide publicly accessible collection of music data.[7] It offers access to over 600,000 CDs and 300,000 LPs, described using international library standards, which it is matching to domain-relevant repositories, including MusicBrainz, Wikidata, sheet music archives and streaming services. The archive, including digitized audio data that can be used for audio analysis and high-quality metadata, makes Muziekweb a relevant authority for classical music in the TROMPA project.

Each of these repositories provides useful information, but their interconnection is limited. Users of these platforms (and many others available on the Web) are often unaware that other platforms exist. Most repositories use their own vocabulary and description standards, and typically do not integrate complementary information available across collections.

4 DATA INFRASTRUCTURE

The data infrastructure centred around the TROMPA *Contributor Environment* (CE) provides for virtual integration of publicly accessible repositories. The CE is a web-based platform that identified and interlinks items hosted within such repositories, collating the combined contributions of TROMPA participants, partners, end-users, and machine processes, coordinating automated tasks and storing metadata pertaining to TROMPA activities.

The CE is implemented as a graph database adhering to FAIR principles [28], employing established Semantic Web ontologies, and making entities retrievable alongside licensing and provenance information via URIs as RDF (JSON-LD[8]). The graph represents metadata describing music content (e.g., score encodings, images, audiovisual recordings), but does not incorporate the content itself; rather, it refers to content hosted in web-accessible public repositories using URIs. Where TROMPA activities generate derivative materials (e.g., crowd-sourced corrections of score encodings; user-contributed performances), these too are contributed to public repositories and referenced from the CE alongside provenance metadata. A GraphQL[9] endpoint provides query access.

Two systems make direct use of the graph within the CE: a collection of automated software agents, and a system to solicit contributions from users through crowd-sourcing tasks. Software agents can query the CE to find new items and perform tasks (e.g., audio signal analysis to calculate key or rhythm information; performance-to-score alignment; or audio transcription). A publish/subscribe notification system allows agents to request notification of new items as they are added to the CE, allowing them to perform their tasks as soon as items become available. Software clients accessing the CE can likewise request notification once a specified task has completed; or, request that users perform actions (e.g., annotation; verification; correction) on items stored in the CE.

Figure 1 illustrates a workflow supported by the CE. An instrumentalist user (see Section 5) wishes to rehearse a piece. The CE
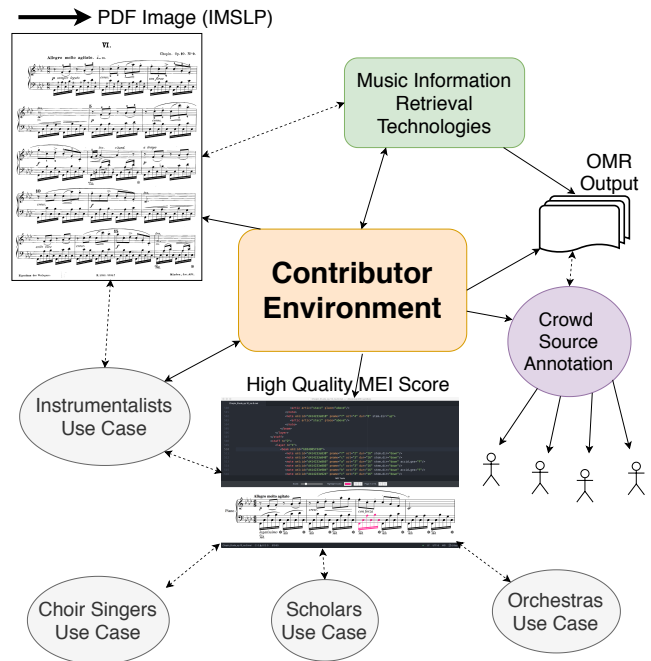

Figure 1: Illustrative PDF-to-MEI conversion schema orchestrated by the TROMPA CE. Solid lines represent direct references; dotted lines are indirect references through the CE.

resolves this to the URI of an IMSLP entry, storing a reference to a PDF of the score hosted there. An OMR task processes the PDF, storing the output at a web-accessible location referenced by the CE. Next, the output (in MEI format) is distributed to crowd annotators in the form of numerous microtasks, selected based on annotator expertise and distributed to maximize the accuracy of crowd improvements with respect to the effort needed. The result of the crowd-sourcing process is a high-quality MEI file[10] stored in a public-domain repository and referenced by the CE. Consequently, the derived MEI score is available not only to the user initiating the process, but to all TROMPA users—and to the wider Web.

5 FIVE USE CASES

TROMPA targets five user types—*music scholars*, *instrumental players*, *choir singers*, *orchestras*, and *music enthusiasts*—through applications providing specialised views of the same underlying knowledge graph. User contributions provide holistic benefits to other users across all use cases. Here, we present an overview of services TROMPA will provide to each type of user, which are informed by user studies conducted over the course of the project. We further consider how each use case benefits from, and provides potential value to, digital musicology.

Scholars. TROMPA allows scholarly users to query interconnected music repositories using the CE, and to contribute insight through scholarly annotation and computer-assisted interpretation (e.g., using feature extraction). All contributions are associated

---

with provenance metadata and furnished with unique identifying URIs, making them citable, reproducible, and reusable in different contexts. Scholars' Web Annotations may target musical scores, audiovisual recordings, digital images, and other digital library resources at various granularities, providing a basis for scholarly discussion and digital multimedia publishing with MELD (per [15]).

Orchestras. TROMPA opens up interesting new possibilities for orchestras, both as performers and as content owners. Renowned conductors interpret the same masterpieces in different ways, reflected by annotated orchestral parts provided to musicians for a given performance project. TROMPA enables orchestra members, librarians, and conductors to efficiently share these written annotations, and to save annotations from conductors' scores for later study. This allows original conducting scores of famous conductors from the past, such as by Willem Mengelberg, to be digitized, their annotations interpreted and linked to MEI score encodings, and audio-visual concert recordings to be aligned, providing rich study material for scholars, musicians, *Kenner und Liebhaber.*

Instrumental players. TROMPA is developing software to support instrumentalists and ensembles in their daily rehearsal regime, enriching rehearsal and teaching situations through immediate feedback on ones own and others' performances [2, 12]. To illustrate: A pianist practicing a novel piece, say Beethoven's *Appassionata*, selects the score on her tablet, retrieving a rendered MEI encoding from the CE. While playing, her performance is streamed as audio or symbolic performance data to an alignment process which synchronizes the performance with the MEI score. After she stops performing, the system displays a note-by-note tempo curve of her performance excerpt. She then selects her favorite performance of the piece by Claudio Arrau from several other performances on YouTube and receives Arrau's tempo curve to compare against her own performance. While listening to Arrau's performance and watching his tempo curve, she writes a comment on a specific section, creating a Web Annotation targeting both the pertinent excerpt of the score and of Arrau's recording, stored with provenance data in the CE. A scholar analysing this particular section during a Beethoven piano sonata project will be able to access her (and other's) annotations of Arrau's performance, incorporate them into his analysis, and reference her comment in his publication.

Choir singers. Choir singers engaging in individual practice will be served by a practicing application enabling them to sing along with synthesized versions of any digital choral score available through the TROMPA data infrastructure, as well as receiving automatically computed feedback on their performance [17] (e.g., regarding intonation and timing). To provide natural singing voice synthesis for choral pieces, we are adapting current models for solo singing based on neural networks [3] using techniques such as voice cloning [4]. Choir singing synthesis will be provided as a service through the Voiceful[11] Cloud API, an existing solution for integrating several speech and singing voice related technologies in a wide variety of applications. Users of the application will generate rich data for exploration by music scholars, revealing those musical excerpts, for instance, that commonly require prolonged practice time; types of intonation error that appear in specific parts; and so

on. At the same time, choir singers can greatly benefit from input by music scholars; for instance, score annotations could be parsed and considered for the synthesis to generate more natural-sounding examples, and relevant annotations could be shown to users of the application to provide greater insight into the piece being practiced.

Music enthusiasts. This use case targets users that lack formal music knowledge but nevertheless enjoy high levels of music engagement. Mediation over arousal and mood through music can have positive effects to support specific daily routines (e.g., learning moments, [5, 24, 25]). TROMPA is developing applications to generate music recommendations aimed to support a mood-boosting or mood-modulating activity, following an audio context-based approach. Music enthusiasts will be able to interact with audio, annotating data related to emotions (i.e., mood, emotions elicited by the musical piece), as well as exploring and discussing annotations provided by other users or software agents. Recommendation systems incorporating advanced classification algorithms with human-in-the-loop feedback provided by the TROMPA infrastructure will generate playlists to help users enhance or change their mood according to their context. Enthusiasts will benefit from scholarly insight in the form of annotations of their favourite music pieces. The open datasets of emotion tags associated with specific pieces and listening contexts will be of interest to music scholars, enabling for instance the correlation of listeners' cultural perspectives and music emotion elicitation. Enthusiasts' discussions are of further interest (per [10]).

6 OPEN CHALLENGES

TROMPA is currently completing its first year of activity, with components largely implemented in prototypical form. To attain the scalability required for the project to be truly useful to digital musicology scholarship, several challenges remain to be addressed. Meaningfully structured, web-addressible music encodings—such as MEI scores—are needed to support fine-grained cross-modal targeting of Web Annotations interconnected with (elements) of musical score, envisioned in TROMPA's use cases, but these remain relatively scarce. We envision that a large community of music lovers of all expertise levels and backgrounds will be able to build and improve a digital library of public-domain music score encodings by applying current MIR technologies with a human-in-the-loop approach as a foundation for enrichment and annotation. Versioning and provenance tracking will serve important roles in crowd motivation. Another challenge lies in the boundless potential for different types of scholarly claims and annotations of interest to digital musicologists. Though the Web Annotation data model supports expression of arbitrary custom motivations, these must be captured and modelled if they are to be expressed and queried consistently by different scholarly users.

To address these challenges, and to serve its overall purpose, it is imperative upon TROMPA to involve the wider digital musicology community in its activities. By adopting an open source, public domain, FAIR, and richly provenanced approach, TROMPA is working towards richer online music public-domain archives for all users.

---

[11]https://www.voiceful.io

ACKNOWLEDGEMENTS

REFERENCES

[1] Andreas Arzt, Harald Frostel, Thassilo Gadermaier, Martin Gasser, Marten Grachten, and Gerhard Widmer. 2015. Artificial Intelligence in the Concertgebouw. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Buenos Aires, Argentina.

[2] Andreas Arzt, Gerhard Widmer, Sebastian Böck, Reinhard Sonnleitner, and Harald Frostel. 2012. Towards a Complete Classical Music Companion. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012),*. Montpellier, France.

[3] Merlijn Blaauw and Jordi Bonada. 2017. A Neural Parametric Singing Synthesizer Modeling Timbre and Expression from Natural Songs. *Applied Sciences* 7, 12 (2017). https://doi.org/10.3390/app7121313

[4] Merlijn Blaauw, Jordi Bonada, and Ryunosuke Daido. 2019. Data Efficient Voice Cloning for Neural Singing Synthesis. *CoRR* abs/1902.07292 (2019). arXiv:1902.07292 http://arxiv.org/abs/1902.07292

[5] Gianna Cassidy and Raymond MacDonald. 2009. The effects of music choice on task performance: A study of the impact of self-selected and experimenter-selected music on driving game performance and experience. *Musicae Scientiae* (2009). https://doi.org/10.1177/102986490901300207

[6] Tim Crawford, Ben Fields, David Lewis, and Kevin Page. 2014. Explorations in Linked Data practice for early music corpora. In *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries*. IEEE Press, 309–312.

[7] David De Roure, Graham Klyne, Kevin R. Page, John P. N. Pybus, and David M. Weigl. 2015. Music and science: Parallels in production. In *Proceedings of the 2nd International Workshop on Digital Libraries for Musicology*. ACM, 17–20.

[8] David De Roure, Graham Klyne, John Pybus, David M. Weigl, and Kevin R. Page. 2018. Music SOFA: An architecture for semantically informed recomposition of Digital Music Objects. In *Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music*. ACM, 33–41.

[9] Ben Fields, Sam Phippen, and Brad Cohen. 2015. A Case Study in Pragmatism: exploring the practical failure modes of Linked Data as applied to classical music catalogues. In *Proceedings of the 2nd International Workshop on Digital Libraries for Musicology*. ACM, 21–24.

[10] Ben Fields and Christophe Rhodes. 2016. Listen To Me–Don't Listen To Me: What Communities of Critics Tell Us About Music.. In *Proceedings of the 17th International Society for Music Information Retrieval Conference*. 199–205.

[11] Martin Gasser, Andreas Arzt, Thassilo Gadermaier, Maarten Grachten, and Gerhard Widmer. 2015. Classical Music on the Web - User Interfaces and Data Representations. In *Proceedings of the 16th International Society for Music Information Retrieval Conference, ISMIR 2015, Málaga, Spain, October 26–30, 2015*. http://ismir2015.uma.es/articles/123_Paper.pdf

[12] Werner Goebl and Gerhard Widmer. 2006. Unobtrusive practice tools for pianists. In *Proceedings of the 9th International Conference on Music Perception and Cognition (ICMPC9)*, Mario Baroni, A. R. Addessi, R. Caterina, and Marco Costa (Eds.). Alma Mater Studiorum University of Bologna, Bologna, Italy, 209–214.

[13] Emilia Gómez, Maarten Grachten, Alan Hanjalic, Jordi Janer, Sergi Jordà, Carles F. Julià, Cynthia C. S. Liem, Agustín Martorell, Markus Schedl, and Gerhard Widmer. 2013. PHENICX: Performances as Highly Enriched aNd Interactive Concert Experiences. In *SMAC Stockholm Music Acoustics Conference 2013 and SMC Sound and Music Computing Conference 2013*. Stockholm, Sweden.

[14] Maria Kallionpää, Chris Greenhalgh, Adrian Hazzard, David M. Weigl, Kevin R. Page, and Steve Benford. 2017. Composing and realising a game-like performance for disklavier and electronics. New Interfaces for Musical Expression.

[15] David Lewis, David M Weigl, Joanna Bullivant, and Kevin R Page. 2018. Publishing musicology using multimedia digital libraries: creating interactive articles through a framework for linked data and MEI. In *Proceedings of the 5th International Conference on Digital Libraries for Musicology*. ACM, 21–25.

[16] Cynthia C. S. Liem, Ron van der Sterren, Marcel van Tilburg, Álvaro Sarasúa, Juan J. Bosch, Jordi Janer, Mark S. Melenhorst, Emilia Gómez, and Alan Hanjalic. 2013. Innovating the Classical Music Experience in the PHENICX Project: Use Cases and Initial User Feedback. In *1st International Workshop on Interactive Content Consumption (WSICC) at EuroITV 2013*. Como, Italy.

[17] Oscar Mayor, Jordi Bonada, and Alex Loscos. 2006. The singing tutor: Expression categorization and segmentation of the singing voice. In *Proceedings of the AES 121st Convention*.

[18] Mark S. Melenhorst and Cynthia C. S. Liem. 2015. Put the concert attendee in the spotlight: A user-centered design and development approach for classical concert applications. In *Proceedings of the 16th International Society for Music Information Retrieval Conference, ISMIR 2015, Malaga, Spain, October 26–20, 2015. Eds. Muller, M., Wiering, F.* IEEE.

[19] Mark S. Melenhorst, Ron van der Sterren, Andreas Arzt, Agustín Martorell, and Cynthia C. S. Liem. 2015. A Tablet App to Enrich the Live and Post-Live Experience of Classical Concerts. In *Proceedings of the 3rd International Workshop on Interactive Content Consumption (WSICC) at TVX 2015*.

[20] Terhi Nurmikko-Fuller and Kevin R. Page. 2016. A linked research network that is Transforming Musicology. In *Proceedings of the 1st Workshop on Humanities in the Semantic Web co-located with 13th ESWC Conference 2016*. 73–78.

[21] Terhi Nurmikko-Fuller, David M. Weigl, and Kevin R. Page. 2015. On organising multimedia performance corpora for musicological study using Linked Data. In *Proceedings of the 2nd International Workshop on Digital Libraries for Musicology*. ACM, 25–28.

[22] Mark Sandler, David De Roure, Steven Benford, and Kevin Page. 2019. Semantic web technology for new experiences throughout the music production-consumption chain. In *2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*. IEEE, 49–55.

[23] Álvaro Sarasúa, Julián Urbano, and Emilia Gómez. 2019. Mapping by Observation: Building a User-Tailored Conducting System From Spontaneous Movements. *Frontiers in Digital Humanities* 6 (2019), 3. https://doi.org/10.3389/fdigh.2019.00003

[24] William Forde Thompson, E. Glenn Schellenberg, and Gabriela Husain. 2001. Arousal, mood, and the Mozart effect. *Psychological Science* (2001). https://doi.org/10.1111/1467-9280.00345

[25] William Forde Thompson, E. Glenn Schellenberg, and Adriana Katharine Letnic. 2012. Fast and loud background music disrupts reading comprehension. *Psychology of Music* (2012). https://doi.org/10.1177/0305735611400173

[26] David M. Weigl, David Lewis, Tim Crawford, Ian Knopke, and Kevin R. Page. 2019. On providing semantic alignment and unified access to music library metadata. *International Journal on Digital Libraries* 20, 1 (2019), 25–47.

[27] David M. Weigl and Kevin R. Page. 2017. A framework for distributed semantic annotation of musical score:" Take it to the bridge!". In *Proceedings of the 18th International Society for Music Information Retrieval Conference*.

[28] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data* 3 (2016).