

# Read/Write Digital Libraries for Musicology

David M. Weigl, Werner Goebel, Alex Hofmann  
{weigl,goebel,hofmann-alex}@mdw.ac.at  
University of Music and Performing Arts Vienna  
Vienna, Austria

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

Tim Crawford, Federico Zubani  
{t.crawford,fzuba001}@gold.ac.uk  
Goldsmiths, University of London  
London, United Kingdom

Alastair Porter  
alastair.porter@upf.edu  
Universitat Pompeu Fabra  
Barcelona, Spain

## ABSTRACT

The Web and other digital technologies have democratised music creation, reception, and analysis, putting music in the hands, ears, and minds of billions of users. Music digital libraries typically focus on an essential subset of this deluge—commercial and academic publications, and historical materials—but neglect to incorporate contributions by scholars, performers, and enthusiasts, such as annotations or performed interpretations of these artifacts, despite their potential utility for many types of users.

In this paper we consider means by which digital libraries for musicology may incorporate such contributions into their collections, adhering to principles of FAIR data management and respecting contributor rights as outlined in the EU’s General Data Protection Regulation. We present an overview of centralised and decentralised approaches to this problem, and propose hybrid solutions in which contributions reside in a) user-controlled personal online datastores, b) decentralised file storage, and c) are published and aggregated into digital library collections. We outline the implementation of these ideas using Solid, a Web decentralisation project building on W3C standard technologies to facilitate publication and control over Linked Data. We demonstrate the feasibility of this approach by implementing prototypes supporting two types of contribution: Web Annotations describing or analysing musical elements in score encodings and music recordings; and, music performances and associated metadata supporting performance analyses across many renditions of a given piece. Finally, we situate these ideas within a wider conception of enriched, decentralised, and interconnected online music repositories.

## CCS CONCEPTS

• **Applied computing** → **Digital libraries and archives**; • **Information systems** → *Multimedia content creation*.

## KEYWORDS

Music Digital Libraries, contribution, decentralisation, Solid, IPFS



This work is licensed under a Creative Commons Attribution International 4.0 License.  
Link to use: <https://creativecommons.org/licenses/by/4.0/>

DLfM '20, October 16, 2020, Montréal, QC, Canada  
ACM ISBN 978-1-4503-8760-6/20/10.  
<https://doi.org/10.1145/3424911.3425519>

## ACM Reference Format:

David M. Weigl, Werner Goebel, Alex Hofmann, Tim Crawford, Federico Zubani, Cynthia C. S. Liem, and Alastair Porter. 2020. Read/Write Digital Libraries for Musicology. In *7th International Conference on Digital Libraries for Musicology (DLfM '20)*, October 16, 2020, Montréal, QC, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3424911.3425519>

## 1 INTRODUCTION

Digital libraries (DLs) take the form of focused collections of (textual or multimedia) digital objects, alongside methods for their access and retrieval, selection, organisation, and maintenance. These may incorporate digital materials interpreting, presenting, describing, or deriving from objects at the core of the collections—tags, annotations, performances, explanations, analyses—without overstretching this definition [4], but such ‘introspective’ materials present architectural challenges. Layered DLs [15], which lift materials describing core collection objects into differentiated layers of presentation, interpretation, or mark-up, elegantly address these challenges. Persistent identifier systems are a critical requirement to making such architectures ‘overlayable, reusable, and repurposable’. Linked Data approaches employing URIs to persistently identify both digital objects and their interrelationships are particularly helpful in implementing layered DLs [22]. By incorporating provenance information, layered DLs accommodate information provided by different agents, potentially with differing authorities.

Such layers may usefully house automatically-generated interpretations and analytical descriptions—so-called ‘feature data’—derived from textual [3] or multimedia [22] digital objects at lower levels of the collection, be they at its core, or aggregations or descriptions at preceding layers. Content for such layers may also be provided by humans. Contributions by interested individuals—music scholars, performers, enthusiasts—are especially useful where subjective aspects, such as overall similarity, quality, or other experiential aspects are to be assessed; to arbitrate cases of ambiguity (e.g., in cases of contested attribution and similar scholarly dispute); or, where contributions take the form of derivative works, such as musical performance recordings or scholarly annotations of texts or scores. Though Web technologies facilitate the elicitation and capture of human perspectives and insights on DL resources, archiving and providing access to metadata and derivative works generated through user contributions tends to be neglected in practice. Co-archival of these perspectives requires DL architectures to explicitly incorporate support for user contributions.

Music DLs are particularly well situated to exercise research into—and benefit from—the structured incorporation of contributed materials into their collections: both owing to their core focus on textual as well as multimedia objects, alongside the reasonable expectation of user contributions in various different modalities; and, to their serving thriving research communities dealing with the derivation of semantic structures from undifferentiated signals, both through machine agents (in the field of music information retrieval) and through human insight (in the fields of empirical musicology and music perception and cognition). A number of music DL projects have explored the utility of user contributions in the form of musicological annotations [2, 14], music performances [18, 32], and accounts of listening experiences [1].

Beside providing descriptions, interpretations, analyses, and other useful content, aggregations of such materials and of their interrelationships with core collection objects and other contributions are themselves interesting targets for analyses. Aggregate analyses of scholarly activity in terms of the ‘footprints of scholarship recorded in library resources’ may yield informative insight in the form of ‘scholarship on scholarship’ [12]. Beyond scholarly activity, aggregated analyses of lay contributions by music enthusiasts (as ‘amateur critics’) have provided insightful musicological findings, relating, e.g., to social constructs of genre in music [11]. Collections of music recordings and performance metadata incorporating different musicians’ renditions of the same piece provide fruitful ground for analyses through the lenses of musicology and performance science [7, 25].

We now consider architectural concerns in supporting structured contributions in DLs, outlining a proposed solution with a focus on user-retained data ownership; present prototypes in two digital musicology case studies that demonstrate the proposal’s feasibility; and contextualise these efforts within broader work to interlink and enrich music repositories on the Web.

## 2 READ/WRITE DL ARCHITECTURES

### 2.1 Centralised and decentralised architectures

Though conceived early on as spaces in which users contribute as well as consume data [6, 8], the predominant paradigm of both the Web and DLs has involved a “read-only” mindset in which information is published by few and consumed by many. Where contributions are supported, user data is usually submitted to a centralised repository—be it content shared on social networks, comments ‘under the line’ on news websites, contributions to Wikipedia articles, or posts to discussions on the ACM DL’s Disqus system. This entails a loss of ownership by the contributor, who transmits control over their content to the institution maintaining the centralised repository according to its terms of service. Additionally, platforms typically maintain their own authentication and access control mechanisms. Together this results in the problem of ‘data silos’ where users cannot easily share or reuse their contributions in other relevant application contexts [17]. At the same time, centralised architectures place significant responsibilities on their maintaining institution, who have legal obligations to safeguard the data and respond to certain user requests (under the rights to access, rectification, erasure, data portability) in accordance with the EU’s General Data Protection Regulation [21] (GDPR).

Recent research into decentralised scholarly communication and knowledge dissemination [28, 31] has explored alternatives to this situation in which contributions are not transmitted into a central repository, but reside in contributor-controlled spaces—*Personal Online Datastores*, or *Pods*—as defined by the Solid project<sup>1</sup> [17]. Beside providing data storage, Solid Pods also act as WebID [26] identity providers, allowing users to authenticate against their own Pod, thus retaining one identity and set of credentials across applications. Instead of ingesting their data, applications may include users’ contributions by reference to URIs within the users’ Pods, using the affordances of Linked Data. Pods are hosted by dedicated services (*Pod providers*) employing standard Web technologies. Open source implementations are available, allowing technically advanced users to host their own Pod servers if desired.

With this approach, users may send their contributions to multiple different applications, allowing for their re-contextualisation and reuse in different contexts in accordance with the FAIR principles of data management [33]. At the same time, users retain ownership and control over their data, including the abilities to keep specific resources private, delegate read and/or write access to specified other users (via reference to their WebID identity), to open access to the public, and of course to delete data where desired.

Beyond empowering the user, this approach also decreases the burden of responsibility faced by DLs and other platforms wishing to incorporate user contributions, as the means to guarantee the rights specified by the GDPR are necessarily retained by the user.

### 2.2 Publishing decentralised contributions

Fully decentralised solutions also have their drawbacks. Users authoring annotations, analyses, or derivative works based upon items in DL collections may wish for these contributions to be discoverable by others. The institutions maintaining DLs may wish to guarantee the persistence and availability of resources integrated with their collections; and may also wish to centralise contributions for reasons of processing efficiency, e.g., if they are to be interpreted in further derivative layers of computational analysis.

There is space between the two extremes of data silos and full decentralisation for hybrid solutions (Figure 1) that achieve a balance between retention of user control over contributed data afforded by Solid, and guarantees of stability and accessibility provided by centralised architectures. We propose an approach in which applications allow users to interact with objects in core DL collections in a decentralised fashion, posting the generated data to users’ own Pods (and potentially sharing access with specified other users or even with the wider public), up to a point of explicit, licensed publication to the DL. Before this point of publication, the generated data exist as Web resources (in a user’s Pod) that reference other Web resources (in the DL’s collection) in the usual fashion of Linked Data. Only after publication do the generated data formally become *contributions* to the DL. There are two possible means in which publication may take place, each striking a slightly different balance between the extremes of centralisation and decentralisation.

- *announce*: URI *references* to contribution copied into the DL;
- *ingest*: contribution *content* copied into the DL.

<sup>1</sup><https://solidproject.org>

Both models place user content within contribution layers of the DL, enriching the DL's collection, and making the content discoverable by its users. In the *announce* model (employed by DL1 in Figure 1), the user retains full control over the data, including over its modification, deletion, or revocation of public access. From the DL's perspective, this minimises overhead around the storage and management of contributions (including GDPR liabilities), but it does require applications interfacing with the DL to be resilient against missing or inaccessible user data. In the *ingest* model (DL2 in the figure), the DL asserts a greater degree of ownership over the contribution, and thus is better able to guarantee that ingested contributions meet data management best practices—e.g., the FAIR principles of *accessibility* (by ensuring data persistence) and *reusability* (by ensuring that contributions are licensed appropriately).

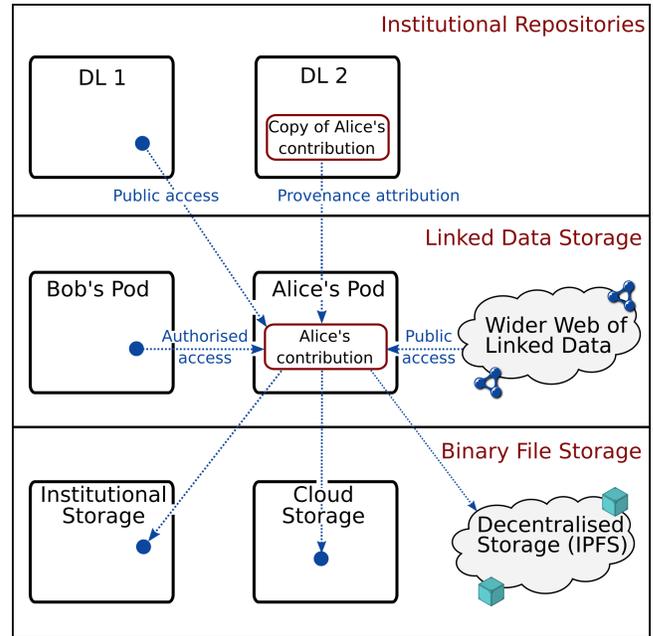
### 2.3 Undifferentiated binary file storage

The Linked Data Platform [19] containers providing file storage and retrieval within Solid Pods are equipped to handle semantically undifferentiated binary files of common interest in digital musicology projects—e.g., audio-visual recordings or high-resolution scanned images—alongside contributions structured as Linked Data. However, storing such files within Solid Pods can be unwieldy, as currently-established Pod providers do not prioritise the provision of large storage quotas or particularly high-bandwidth access.

As Web resources, large binary files can be referenced by Linked Data residing within a user's Pod, even when the files are stored elsewhere. One option is to open up a DL's centralised institutional storage to submissions of such data, but this places significant administrative and networking overhead upon the institution, situates contributions firmly within the 'silo' of the particular DL, and diminishes contributors' control over their data (Section 2.1). An alternative is to direct contributions of such files to cloud-hosting services such as those provided by Amazon S3 (see e.g., [29]); files stored in such locations may be referenced from different repositories but data ownership would still be transferred to the institution holding the access credentials, unless users are constrained to operating their own S3 accounts, a non-trivial requirement. Further, the data's availability still hinges on the single cloud service provider.

An alternative to file storage with a centralised provider are decentralised peer-to-peer networks in which users publish and share their files themselves. The InterPlanetary File System<sup>2</sup> (IPFS) is such a network, incorporating a content-addressed data model for file storage and access with its own corresponding protocol [5, 30]. The network consists of connected nodes that transfer files and other data structures. Unlike HTTP, files in IPFS are addressed using *content addressing*, where each file uploaded to the network is identified by a unique multihash checksum, a short machine-readable fingerprint of the file content. This reduces dependency on specific centralised servers, as multiple nodes may hold copies of the same file. Latency is also minimised as files can be accessed at the closest nodes in the network. Users may serve their undifferentiated binary files via IPFS and publish the hash links via their Solid Pods [24]. These links can be resolved directly using an IPFS-compliant client or by using an HTTP gateway.

<sup>2</sup><https://ipfs.io>



**Figure 1: Layered contribution architecture.** Dashed lines: URI references. Contributions in Pods (Linked Data storage) are published to institutional repositories (Digital Libraries—DLs) by *announce* (DL1) or *ingest* (DL2). Large binary files are stored with a centralised provider, or in decentralised fashion (e.g., by InterPlanetary File System—IPFS).

## 3 CASE STUDIES IN DIGITAL MUSICOLOGY

TROMPA [16] is a project to interconnect public-domain music resources published in repositories on the Web, and to enrich these resources using machine processes and user contributions. Within this project we have implemented prototypical applications for two use-cases—respectively focusing on annotation of scores and audio recordings, and on the analysis and review of rehearsal attempts—that together serve to demonstrate the feasibility of the ideas presented in this paper. In each case, contributions are generated using Web applications that store data in users' Solid Pods, from where they may be published to TROMPA's data infrastructure.

### 3.1 Web Annotation of music resources

Published musicology (at any level of sophistication) is traditionally presented as text, accompanied by scores, illustrations and diagrams necessary to communicate the author's argument; these are essentially static and are necessarily fixed on the printed page. Digital media offer a vast range of possibilities of extending the form, style, scope and range of such presentations, one example being the inversion of this process and providing a complete score as entry point which is annotated as required by links, labelled on the musical staves, to textual commentary, or to other music examples (extracts or complete works) in various formats, e.g., static graphical diagrams or scores, or audio/video recordings, as desired. While this is directly analogous to the now familiar notion of hypertext, technical restrictions due to the complexities of music have hitherto

limited the extent to which this has been attempted—though some recent publications of musicological scholarship [10, 13] achieve this using the Music Encoding and Linked Data (MELD) framework.

We have implemented a set of Javascript/React modules facilitating the creation of annotation tools for such materials published on the Web. The TROMPA annotation module for Music Encoding Initiative (MEI) score encodings<sup>3</sup>, a MELD wrapper component, supports click-and-drag selection of musical elements of configurable granularity (e.g., notes; measures), and authentication and submission of annotations to a Solid Pod. The audio annotation module<sup>4</sup> provides a tool for annotation and playback of audio signals, visualised using *Wavesurfer.js*<sup>5</sup>. We are developing user interfaces combining these modules,<sup>6</sup> informed by requirement-gathering and iterative validation through a series of user studies.

The Web Annotations data model [27] provides a flexible, standardised way of expressing annotations about (fragments of) Web resources. Our tooling generates Web Annotations targeting MEI fragments according to the identifiers of user-selected score elements, and media fragments of audio recordings representing instants and intervals at timed offsets from the beginning of the recording. The Web Annotation standard specifies a set of 13 default annotation motivations identified in a survey of the annotation landscape—including describing, highlighting, and linking—which encode the annotation’s purpose in machine-readable form, and may thus be used in rendering or interaction logic by user interfaces. These are easily extended with custom motivations; indeed, we are exploring the creation, publication, and reuse of custom annotation motivations themselves as another layer of user contribution. One further convenience provided by this standard is the canonical property. Using this property, the annotation’s URI within the contributing user’s Pod is specified as the canonical identifier for the annotation. This neatly simplifies provenance tracking and disambiguation where annotations are duplicated, as is the case where annotations are published to a DL—or conceivably multiple independent repositories—by *ingest* (Section 2.2).

### 3.2 Score-aligned performances

Analyses in musicology and performance science may investigate multiple recorded renditions of the same piece, potentially performed by different individuals, at different times, or in different contexts [7, 25]. Such investigations benefit from the fine-grained alignment of recorded audio signals and timed performance metadata—signal-derived features such as note-onset and -offset times, timbral or dynamic descriptors, or instrument-specified features supplied by digital instruments (e.g., MIDI events, key trajectories, or pedal positions on a Disklavier or Bösendorfer SE piano)—with encoded music score elements [9]. The Music and Timeline Ontologies [23] offer convenient Linked Data structures conceptualising these metadata associations as instants or intervals specified along a performance timeline, potentially corresponding to a recorded audio signal. Such timelines may be aligned to encoded scores, either through manual expert annotation [14] or by automated techniques (e.g., [20]); where such alignments are in place and corresponding

audio recordings available, the comprehensive capture of musical elements within score encodings supply a semantic index into undifferentiated audio files recording the performance signal [14].

Within TROMPA we are developing a prototype application primarily targeting music performers<sup>7</sup> that realises these ideas. The application enables musicians to record and score-align their rehearsal attempts, offering score-informed review and analysis of their rehearsal progress over time. Linked Data structures representing the timelines are generated for each recorded rehearsal as described above, and stored within the performer’s Solid Pod. For each thus recorded rehearsal, performers may take advantage of the affordances described in Section 2.1 to retain private access to the corresponding timeline and associated metadata; to share them with specified other users via their WebID; or, to publish them under public license to the TROMPA data infrastructure, where they become discoverable and accessible to a wider community of performers, scholars, and enthusiasts [32]. Where large numbers of renditions, potentially from many different contributors, are available, this enables the pursuit of interesting and novel musicological investigations: for instance, establishing quantitative measures of the degree of expressive variation across the corpus of interpretations of a given piece; or, empirical measures of score difficulty, informed by the distribution of performance errors (inserted or omitted notes, as identified during the score alignment process) across renditions and across individual performers.

## 4 CONCLUSION

Though conceived early on as read/write platforms for information interchange, DLs—and the wider Web—have predominantly operated in a publish/consume paradigm, underemphasising decentralised contribution of information. In this paper, we have outlined the value of supporting contributions of ‘introspective’ materials describing or deriving from objects within core DL collections. We have discussed the architectural and legislative challenges of incorporating publication and storage of such contributions in layered DLs, and have presented two case studies contextualising our proposed hybrid approach combining decentralised personal online datastores and centralised publication to institutional repositories within use cases relevant to digital musicology.

As a publicly funded project in this area, TROMPA has a strong focus on EU data governance-compliant data handling. Modern Web (re-)decentralisation projects such as Solid are working toward more equitable processes of information exchange. TROMPA’s use cases [16] outline a spectrum of demands regarding contribution, access, and governance of user data. This provides a fruitful setting to consider not just architectural concerns, but also the user-facing reception and implications of such solutions to a varied selection of audiences, including those with limited technical expertise; and corresponding challenges of eliciting and incentivising contributions, and providing for their reuse in different contexts, potentially beyond those foreseen within project scope. Through ongoing research and development in this area, we hope to establish new insights into how music DLs may more fully embrace—and benefit from—user-contributed creativity, expertise, and insight relating to or deriving from the materials within their collections.

<sup>3</sup><https://github.com/trompamusic/selectable-score>

<sup>4</sup><https://github.com/trompamusic/audio-annotator>

<sup>5</sup><https://wavesurfer-js.org>

<sup>6</sup><https://github.com/trompamusic/music-scholars-annotator>

<sup>7</sup><https://github.com/trompamusic/clara>

## ACKNOWLEDGMENTS

This project has received funding from 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 collaboration of our colleagues in the TROMPA consortium.

## REFERENCES

- [1] Alessandro Adamou, Simon Brown, Helen Barlow, Carlo Allocca, and Mathieu d'Aquin. 2019. Crowdsourcing Linked Data on listening experiences through reuse and enhancement of library data. *International Journal on Digital Libraries* 20, 1 (2019), 61–79.
- [2] David Bainbridge. 2015. And we did it our way: A case for crowdsourcing in a digital library for musicology. In *Proceedings of the 2nd International Workshop on Digital Libraries for Musicology*. 1–8.
- [3] David Bainbridge, J. Stephen Downie, and Boris Capitanu. 2018. Providing pin-point page-level precision to 1 trillion tokens of text for worksret creation. In *Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries*. 407–408.
- [4] David Bainbridge, Xiao Hu, and J. Stephen Downie. 2014. A musical progression with Greenstone: How music content analysis and Linked Data is helping redefine the boundaries to a music digital library. In *Proceedings of the 1st International Workshop on Digital Libraries for Musicology* (London, United Kingdom) (DLfM '14). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/2660168.2660170>
- [5] Juan Benet. 2014. IPFS-content addressed, versioned, P2P file system (DRAFT 3). <https://ipfs.io/ipfs/QmR7GSQM93Cx5eAg6a6yRzNde1FQv7uL6X1o4k7zrJa3LX/ipfs.draft3.pdf>
- [6] Tim Berners-Lee and Kieron O'Hara. 2013. The read–write linked data web. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371, 1987 (2013), 20120513.
- [7] Nicholas Cook. 2014. Between art and science: Music as performance. *Journal of the British Academy* 2 (2014), 1–25.
- [8] Gregory Crane, David Bamman, Lisa Cerrato, Alison Jones, David Mimno, Adrian Packel, David Sculley, and Gabriel Weaver. 2006. Beyond digital incunabula: Modeling the next generation of digital libraries. In *Proceedings of the International Conference on Theory and Practice of Digital Libraries*. Springer, 353–366.
- [9] Johanna Devaney and Hubert Léveillé Gauvin. 2016. Representing and linking music performance data with score information. In *Proceedings of the 3rd International workshop on Digital Libraries for Musicology*. 1–8.
- [10] Laurence Dreyfus, David Lewis, and Kevin R. Page. 2020. An interactive multimedia companion to Wagner's Lohengrin: encoding and visualising a motivic study. In *DH 2020*. <https://hcommons.org/deposits/item/hc:31987>
- [11] 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.
- [12] David Lewis, Yun Fan, Glenn Henshaw, and Kevin R. Page. 2017. Musicology of digital libraries: structure in RILM. In *Proceedings of the 4th International Workshop on Digital Libraries for Musicology*. 59–62.
- [13] 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*. 21–25.
- [14] David Lewis, David M. Weigl, and Kevin R. Page. 2019. Musicological observations during rehearsal and performance: A Linked Data digital library for annotations. In *Proceedings of the 6th International Conference on Digital Libraries for Musicology*. 1–8.
- [15] Clifford Lynch. 2002. Digital collections, digital libraries and the digitization of cultural heritage information. *First Monday* 7, 5 (2002).
- [16] David M. Weigl, Werner Goebel, Tim Crawford, Aggelos Gkiokas, Nicolas F. Gutierrez, Alastair Porter, Patricia Santos, Casper Karreman, Ingmar Vroomen, Cynthia CS Liem, et al. 2019. Interweaving and enriching digital music collections for scholarship, performance, and enjoyment. In *Proceedings of the 6th International Conference on Digital Libraries for Musicology*. 84–88.
- [17] Essam Mansour, Andrei Vlad Sambra, Sandro Hawke, Maged Zereba, Sarven Capadislis, Abdurrahman Ghanem, Ashraf Aboulnaga, and Tim Berners-Lee. 2016. A demonstration of the solid platform for social web applications. In *Proceedings of the 25th International Conference Companion on World Wide Web*. 223–226.
- [18] Albert Meroño Peñuela, Reinier de Valk, Enrico Daga, Marilena Daquino, and Anna Kent-Muller. 2018. The Semantic Web MIDI tape: An interface for interlinking MIDI and context metadata. In *Proceedings of the 1st International Workshop on Semantic Applications for Audio and Music* (Monterey, CA, USA) (SAAM '18). Association for Computing Machinery, New York, NY, USA, 24–32. <https://doi.org/10.1145/3243907.3243909>
- [19] Nandana Mihindukulasooriya and Roger Menday. 2015. *Linked Data Platform 1.0 primer*. Technical Report. <http://www.w3.org/TR/ldp-primer/>
- [20] Eita Nakamura, Kazuyoshi Yoshii, and Haruhiro Katayose. 2017. Performance error detection and post-processing for fast and accurate symbolic music alignment. In *Proceedings of the 18th International Society for Music Information Retrieval Conference*. 347–353.
- [21] Official Journal of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). <https://op.europa.eu/en/publication-detail/-/publication/3e485e15-11bd-11e6-ba9a-01aa75ed71a1>
- [22] Kevin R. Page, Sean Bechhofer, Gyorgy Fazekas, David M. Weigl, and Thomas Wilmering. 2017. Realising a layered digital library: exploration and analysis of the live music archive through linked data. In *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. IEEE, 1–10.
- [23] Yves Raimond, Samer A. Abdallah, Mark B. Sandler, and Frederick Giasson. 2007. The Music Ontology. In *Proceedings of the 8th International Conference on Music Information Retrieval*.
- [24] Manoharan Ramachandran, Niaz Chowdhury, Allan Third, John Domingue, Kevin Quick, and Michelle Bachler. 2020. Towards complete decentralised verification of data with confidentiality: Different ways to connect Solid Pods and Blockchain. In *Companion Proceedings of the Web Conference 2020*. 645–649.
- [25] Bruno H. Repp. 1992. Diversity and commonality in music performance: An analysis of timing microstructure in Schumann's "Träumerei". *The Journal of the Acoustical Society of America* 92, 5 (1992), 2546–2568.
- [26] Andrei Sambra, Henry Story, and Tim Berners-Lee. 2014. *WebID Specification*. Technical Report. <http://www.w3.org/2005/Incubator/webid/spec/identity/>
- [27] Robert Sanderson, Paolo Ciccarese, and Benjamin Young. 2017. *Web Annotation Data Model (W3C Recommendation)*. Technical Report. <https://www.w3.org/TR/annotation-model/>
- [28] Sarven Capadislis. 2020. *Linked research on the decentralised Web*. Ph.D. Dissertation. Rheinische Friedrich-Wilhelms-Universität Bonn. <https://csarven.ca/linked-research-decentralised-web>
- [29] Victor Jesus Sosa-Sosa and Emigdio M. Hernandez-Ramirez. 2012. A file storage service on a cloud computing environment for digital libraries. *Information Technology and Libraries* 31, 4 (2012), 34–45.
- [30] Mathis Steichen, Beltran Fiz, Robert Norvill, Wazen Shbair, and Radu State. 2018. Blockchain-based, decentralized access control for IPFS. In *2018 IEEE International Conference on Internet of Things (IThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*. IEEE, 1499–1506.
- [31] Ruben Verborgh. 2020. Re-decentralizing the Web, for good this time. In *Linking the World's Information: Tim Berners-Lee's Invention of the World Wide Web*, Oshani Seneviratne and James Hendler (Eds.). ACM. <https://ruben.verborgh.org/articles/redecentralizing-the-web/>
- [32] David M. Weigl and Werner Goebel. 2020. Rehearsal encodings with a social life. In *Music Encoding Conference Proceedings 2020*. <https://hcommons.org/deposits/item/hc:31983>
- [33] 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, 1 (2016), 1–9.