Collaborative annotation and semantic enrichment of 3D media

A FOSS toolchain

Lozana Rossenova
Open Science Lab
TIB – Leibniz Information Centre
for Science and Technology
Hannover, Germany
lozana.rossenova@tib.eu

Zoe Schubert
SPK – Stiftung Preussischer
Kulturbesitz
Berlin, Germany
Zoe.Schubert@sbb.spk-berlin.de

Richard Vock
Institute of Computer Science
University of Bonn
Bonn, Germany
vock@cs.uni-bonn.de

Lucia Sohmen
Open Science Lab
TIB – Leibniz Information Centre
for Science and Technology
Hannover, Germany
lucia.sohmen@tib.eu

Lukas Günther
Open Science Lab
TIB – Leibniz Information Centre
for Science and Technology
Hannover, Germany
lukas.guenther@tib.eu

Paul Duchesne
Open Science Lab
TIB – Leibniz Information Centre
for Science and Technology
Hannover, Germany
paul.duchesne@tib.eu

Ina Blümel
Open Science Lab
TIB – Leibniz Information Centre
for Science and Technology
Hannover, Germany
ina.blumel@tib.eu

ABSTRACT

A new FOSS (free and open source software) toolchain and associated workflow is being developed in the context of NFDI4Culture, a German consortium of research- and cultural heritage institutions working towards a shared infrastructure for research data that meets the needs of 21st century data creators, maintainers and end users across the broad spectrum of the digital libraries and archives field, and the digital humanities. This short paper and demo present how the integrated toolchain connects: 1) OpenRefine – for data reconciliation and batch upload; 2) Wikibase – for linked open data (LOD) storage; and 3) Kompakkt – for rendering and annotating 3D models. The presentation is aimed at librarians, digital curators and data managers interested in learning how to manage research datasets containing 3D media, and how to make them available within an open data environment with 3D-rendering and collaborative annotation features.

CCS CONCEPTS

• Data management systems • Open Source Software

KEYWORDS

OpenRefine, Wikibase, Kompakkt, Linked Open Data, 3D data

1 Introduction

Mass-digitisation efforts have resulted in the launch of numerous online archives and collections [1]. Libraries and other cultural heritage institutions have had to adapt their practices to account for shifts in what digital cultural stewardship entails [2]. Increased accessibility and reuse of data is key [3; 4], but in the vast majority of cases, access implies viewing siloed resources, such as scanned text pages or photographs of physical objects and sometimes 3D renderings, alongside minimal descriptive metadata. But cultural

memory preservation is not guaranteed through digitisation activities alone [5]. Instead, it requires the active participation of diverse audiences who can search, access and enrich datasets through annotation and critical interpretation – the latter being especially important for researchers. Collaborative environments that facilitate easy participation (e.g. wikis) and follow established standards (e.g. IIIF), have already been deployed successfully across the institutional and non-profit sector, but only when it comes to dealing with digitized texts and 2D images. Digital representations of cultural assets in the form of 3D models within disciplines such as architecture, art, archaeology, and 3D reconstruction are particularly heterogeneous in formats and structure [6], ergo standardized access and visualisation tools fail to meet new research objectives and institutional requirements.

3D datasets pose unique challenges to renderers in terms of geometric complexity, memory and bandwidth requirements [7]. Existing tools that tackle these challenges, often do so at the expense of other access priorities – such as rich metadata, collaborative and participatory environments, or standardized approaches to data exchange for e.g. annotations [8]. To address this knowledge gap, we have developed a suite of tools as part of a larger national effort which involves the partnership between several research, library and cultural institutions. In addition, we have placed particular emphasis on collaborative development, which is already well established in FOSS (Free and Open Source Software) communities but rather a novelty in infrastructure projects.

The proposed toolchain focuses on the annotation of 3D data within a collaborative knowledge graph environment, so that 3D objects' geometry, attendant metadata, as well as user-contributed annotations remain searchable, and data interconnections are not lost. The project builds on several existing FOSS tools:

- OpenRefine: a data cleaning, reconciliation and batch upload tool with a graphical user interface [9];

- Wikibase: a suite of services developed by Wikimedia Germany to manage, query and access linked open data within a version-controlled environment [10];

- Kompakkt: a browser-based open-source 3D and multimedia viewer with built-in collaborative annotation features [11].

This suite of tools follows FAIR principles, and adopts common standards like persistent identifiers (PIDs) and the W3C web annotation model. By integrating these state-of-the-art tools, we are extending their individual capabilities: until now it was not possible to work with 3D data in Wikibase, and embedding semantic concepts within media annotations was not possible within Kompakkt. Our toolchain facilitates linking 3D objects and annotations, and their cultural context (including historical people and places, geo-location and 3D-capture-technology metadata), to the broader semantic web and various national and international authority records (GND, Getty's AAT, VIAF and more).

2 Methods

The implementation of the Open Refine, Wikibase and Kompakkt toolchain is carried out by our interdisciplinary team of designers, developers and information scientists in an open agile process in close collaboration with the research and cultural communities affiliated with the project, as well as the open source developer communities around these tools, in order to increase transparency, reduce duplication and ensure optimal usability.

In Phase 1, we developed a minimum viable product (MVP) based on the user stories identified in the course of establishing the original NFDI4Culture infrastructural proposal [12]. Furthermore, in the course of developing the MVP, we used data from a specific architectural case study – reconstruction work and 3D modelling of Weikersheim Castle in Germany and its painted interiors [13]. This allowed us to work with real world data and develop all aspects of the toolchain with concrete user requirements in mind. The case study combines a wide range of heterogeneous research data and associated metadata: from flat floor plans, 3D reconstructions of entire halls, photographs of paintings and interiors, video walkthroughs, to long-form art historical texts, alongside related historical persons and geo locations; in order to serve as proof of concept that this toolchain can be applied beyond the field of architecture in a wide range of related cultural disciplines.

Following the release of the beta version of the MVP in March 2021, we have planned various user testing activities, including workshops and demos, which will seek further community feedback.

After a period of testing and bug fixing, we will begin Phase 2 of the project development, which will focus on improving documentation, extending use-cases and scaling efficient deployment strategies for multiplying toolchain instances across different datasets and different knowledge domains. Open-standards-conformance will once again play an important role here, which will involve community work and engagement with standards bodies, such as the IIIF 3D Technical Specification Group [14].

3 FOSS Toolchain: technical architecture and user-facing interaction model

The toolchain is designed with a modular architecture (Figure 1) wherein all data is uploaded to a Wikibase repository via POST requests and it can be queried via a dedicated SPARQL endpoint. By abstracting the storage layer, we can switch the choice of databases required for both Wikibase and Kompakkt and upgrade to, for example, different triplestore technology in case of future maintenance or obsolescence concerns. Data across all three stages of a data workflow (cleaning and upload, storage and indexing, and presentation and annotation) is readily editable via graphical user interfaces, making the toolchain accessible to a broad public of researchers and data managers, and lowering the learning curve for working with linked open data (LOD). Long-term sustainability, efficient and scalable deployment, as well as

increased usability across the whole suite of tools are the main guiding principles of our development process.
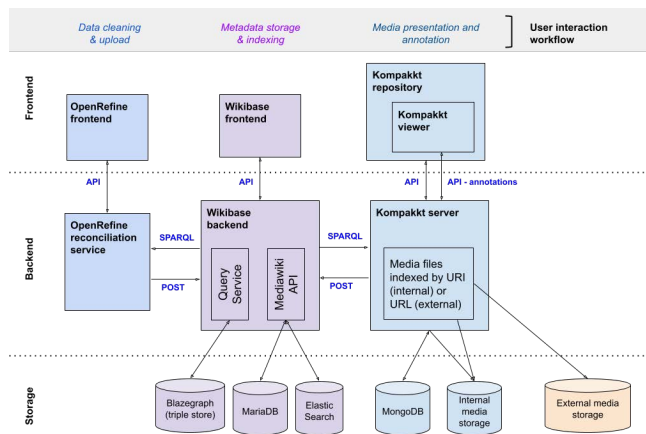

Figure 1: Diagrammatic representation of the FOSS toolchain architecture.

3.1 OpenRefine

Previously Google Refine, OpenRefine is "a powerful tool for working with messy data: cleaning it; transforming it from one format into another; and extending it with web services and external data." [15]. Most crucially for projects intending to use Wikibase and/or Wikidata as a final data repository, OpenRefine allows direct manipulation of data in Wikidata via a reconciliation service and an editing extension. Since version 3.5 (officially released in 2021), it allows users to connect to any Wikibase instance. However, the process of connecting OpenRefine to individual Wikibase instances remains poorly documented and requires a lot of "insider" knowledge accessible only within specific community groups.

The first milestone of the toolchain orchestration was establishing a mechanism for automated deployment of reconciliation services between OpenRefine and an arbitrary, self-managed instance of Wikibase. This was completed via an Ansible playbook and CI/CD pipelines on GitLab [16].

Next, we developed a clearly documented pipeline for cleaning data, reconciling it against concepts from authority control services in order to enrich our test dataset with additional data that could be inferred from the reconciled relations, and finally uploading to a Wikibase instance [17]. It is worth noting that once the data is added to a Wikibase instance, it can also serve as a source for enriching other datasets that could be reconciled against it in OpenRefine, including Wikidata.

3.2 Wikibase

Wikibase and Wikidata are two related software packages from the Wikimedia family of applications. They allow the storage and management of structured machine-readable data following LOD protocols, besides featuring other common characteristics of wikis, such as collaboration and version control features.

Wikidata is the public-facing instance of the software – a Wikipedia for structured data; it is maintained by Wikimedia Germany. Wikibase is the open source software environment built to run Wikidata. Crucially, Wikibase can be deployed independently from Wikidata (and Wikimedia) and can be customized to suit the needs of individual data domains and data repositories [18].

The Wikibase backend connects a standard MySQL database to a triplestore – writing the edits created by users via the frontend interface to RDF in the process. In order to allow data to remain editable, all data changes have to be made through the relational database, and only then streamed to the triplestore and the attendant public SPARQL endpoint. Running this particular architecture on a virtual machine is computationally intensive, and comes with various performance caveats (e.g. slow writing speed for bulk uploads via the API). Deploying multiple instances for various institutional or domain-specific use cases inevitably results in extended maintenance and growing computing resources. Therefore, in addition to developing an efficient deployment mechanism for OpenRefine services, we have conducted early testing with a Kubernetes orchestration for running multiple Wikibase instances – WBStack [19]. This is an open source solution maintained by Wikimedia Germany and an official cloud service is launching in 2022 [20]. Deploying instances via this cloud-based Kubernetes cluster would increase resource efficiency, reduce maintenance requirements, and would still be possible to connect with dedicated OpenRefine reconciliation services.

Still, Wikibase is primarily a data management environment, which is not optimised to render large media files, such as 3D media or AV. Hence the need to extend the toolchain with a dedicated media viewer with annotation capabilities, such as Kompakkt.

3.3 Kompakkt

Once all relevant data is modeled and uploaded in a Wikibase repository, the next milestone is linking the Wikibase database with the rendering and user interaction environment of Kompakkt. Kompakkt offers a modern frontend framework [21] for user interaction and thanks to its own relational database offers users a collaborative environment for annotation 2D, 3D and audio-video media. The metadata layer for description of media, as well as the annotation's content are all following the W3C web annotation standard and are mapped to CIDOC-CRM.

What Kompakkt doesn't offer out-of-the-box is a semantic layer and the possibility to add richer context about media items and their relations, such as – how a room relates to its location within a larger building construction, for instance. With the Wikibase-Kompakkt integration, we are opening up this possibility by mapping Kompakkt's entire data model to a Wikibase repository, and extending it with further details about people, locations, and technical processes connected to the creation of the media files. When uploading a new media file to Kompakkt, users simultaneously generate a new item in Wikibase and all metadata

attached to that media file is then stored in Wikibase. Additionally, during the upload workflow users have access to all the concepts already uploaded to Wikibase, so there is no data redundancy. The only aspect of data writing that remains within the original design of Kompakkt is the user agent data, which is stored in Kompakkt's MongoDB and controls rights for media upload and/or collaborative annotation authoring. User agents are also stored as items in Wikibase, so data provenance is preserved in the semantic storage, too, making it possible to query via the SPARQL endpoint.

By enriching Kompakkt's existing architecture with the semantic environment of Wikibase, we are able to infer more complex relations between media items and their contextual metadata than previously possible. 3D reconstructions of rooms can be linked to their precise location within a building's flat plan, or to handscanned media representations of objects in the room, such as sculptures or ceiling paintings. The latter (i.e. the media representations) can be described with rich art historical data about the physical objects themselves, beyond the narrow technical metadata for the 3D file. Thus we are able to connect knowledge and research data produced by architects, reconstruction technicians and art historians, among others. What is more, all this data can be added and edited collaboratively via user-facing frontend interfaces.

3.4 Frontend user interaction

The toolchain is envisioned as a collaborative environment with different levels of read/write access, wherein groups of researchers can work collaboratively on data upload and annotation with clear data provenance. To allow for various levels of access and maintenance of the research data, different capabilities are provisioned for users logged in via the Wikibase or the Kompakkt interface.

Users with access to Wikibase are able to bulk upload data (via OpenRefine), add to and edit the data model as needed (i.e. controlling classes and properties), and add all core concepts, people, as well as physical objects and locations, that will be then made available for browsing and reuse during annotation in Kompakkt. Wikibase users in this context should be the digital librarians, curators or researchers tasked with core data management roles within a particular project or institutional use-case. There are two further tiers of user access rights management via Kompakkt's interface: 1) *users with upload rights* are able to upload media (2D, 3D and more), attach metadata to media by either browsing and selecting from relevant concepts already existing in the Wikibase triplestore, or by manually entering data such as dates, URLs, or plain text descriptions, and last but not least they are able to create annotations which can combine a mix of manually-entered text and connections to existing items in Wikibase; 2) *users with non-upload rights* can only annotate media objects and save these into personal collections, thus helping to establish a distinction between annotations added by the creator (or uploader) of a media item, and personal annotations added in the context of doing research with an existing media item. Kompakt users with upload rights can be senior researchers and

media creators within a research project, whereas non-upload-rights users can be junior researchers, students, or – in some cases – members of the general public. All annotations created by any user with a Kompakkt account get stored as individual items (with a PID) in Wikibase, remaining searchable via ElasticSearch and the SPARQL endpoint. Search via either Wikibase or Komapkkt's interface remains open to all users – whether logged in, or not.

Note that logins between Wikibase and Kompakkt are distinct, so that Kompakkt users can remain separate from users who also have rights to edit the data on the level of the ontology and data model. At the same time Wikibase administrators will need to create a separate account with upload rights within Kompakkt, if they also wish to contribute media files and annotations. In future development phases of the toolchain, more granular access rights management and integration with single-sign-on systems will be further areas of focus, with the aim of making user interaction with the toolchain better suited to various institutional contexts.

4 Outlook

Knowledge graphs and semantic web solutions have promised to liberate data from institutional and disciplinary silos, but most existing software for research data management remains focused on text-only data. Software focused on the presentation of non-textual media, and particularly 3D models, tends to use 'flat storage' and lacks options for collaboration, version control and accurate provenance tracking. The FOSS toolchain presented in this paper addresses these complementary issues without reinventing the wheel. By taking a modular approach to the data architecture, the data model mapping, and the frontend interface customization, we are proposing to use already working, non-proprietary solutions with a view towards long-term sustainability.

Modularity and flexibility are also intended to keep the toolchain readily applicable to different research and knowledge domain scenarios, which may also feature other data types, such as AV media. At the same time, thanks to the capacity to store data in RDF format and the public SPARQL endpoint of any Wikibase instance, deploying multiple instances of the toolchain for multiple project partners does not mean that their resources will be siloed. On the contrary, the data can remain interoperable within a federated system wherein multiple collections, whose data and metadata are stored in independent instances and follow standards specific to the data provider's domain, can still be queried from a single endpoint. Improving bulk operations and automation features within the semantic structure of annotations (e.g. entity recognition) remain open goals for further development of our toolchain, so that annotated non-textual media can stop being secondary to textual research data and become a usable resource within a federated and sustainable semantic web ecosystem.

ACKNOWLEDGMENTS

REFERENCES

[1] Terras, M. 2011. The rise of digitization. In Rikowski, R. (ed.), *Digitisation Perspectives*. Rotterdam: Sense Publishers, pp. 3–20.

[2] Parry, R. (ed) 2010. *Museums in a Digital Age*. Abingdon: Routledge.

[3] Kapsalis, E. 2016. The impact of open access on galleries, libraries, museums, & archives. Available from: http://siarchives.si.edu/sites/default/files/pdfs/2016_03_10_OpenCollections_Public.pdf (Accessed 15 July 2021)

[4] Maher, K. and Tallon, L. (2018). Wikimedia and The Met: A shared digital vision, *Wikimedia Blog*. https://blog.wikimedia.org/2018/04/19/wikimedia-the-met-shared-digital-vision/ (Accessed 15 July 2021)

[5] Levi, A. S. 2022. Filling the Gaps: Digital Humanities as Restorative Justice. *DHd 2022 – Potsdam, Online, 07-11 March 2022.* Available from: https://www.dhd2022.de/opening-keynote/ (Accessed 12 March 2022)

[6] Blümel, I. and Wessel, R. 2019. DDB goes 3D. *Zenodo.* [online] DOI: 10.5281/zenodo.5579159 (Accessed 07 Jan 2022)

[7] Koller, D., Frischer, B. and Humphreys, G. 2009. Research challenges for digital archives of 3D cultural heritage models, *JOCCH* 2, (7): 10.1145/1658346.1658347.

[8] Fund, N., Schoueri, K. and Scheibler, C. (2021) *Pure 3D: Technical Report*, Maastricht University. Available from: https://pure3d.eu/wp-content/uploads/2021/09/Pure3D_Technical-Report.pdf (Accessed 18 Jan 2022)

[9] See: https://openrefine.org/. See also: Sterner, S.. 2019. Cleaning Collections Data Using OpenRefine. In*: Issues in Science and Technology Librarianship.* 92. DOI: 10.29173/istl30 (Accessed 18 Jan 2022)

[10] See: https://wikiba.se/. See also: Alípio, S., Abdulai, M.S., Burnett, G. and Shick, D. 2021. Wikibase: the Software for Open Data projects. *Wikimedia Tech News.* https://tech-news.wikimedia.de/en/2021/04/14/wikibase-the-software-for-open-data-projects/ (Accessed 18 Jan 2022)

[11] See: https://kompakkt.de/. See also: Eide, Ø., Schubert, Z., Türkoğlu, E., Wieners, J.G. and Niebes, K. 2019. The intangibility of tangible objects: re-telling artefact stories through spatial multimedia annotations and 3D objects. *Presented at the ICOM Kyoto 2019, 25th ICOM General Conference: Museums as Cultural Hubs: The Future of Tradition, Kyoto.* http://doi.org/10.5281/zenodo.3878966 (Accessed 18 Jan 2022)

[12] Altenhöner R, Blümel I, Boehm F, Bove J, Bicher K, Bracht C, Brand O, Dieckmann L, Effinger M, Hagener M, Hammes A, Heller L, Kailus A, Kohle H, Ludwig J, Münzmay A, Pittroff S, Razum M, Röwenstrunk D, Sack H, Simon H, Schmidt D, Schrade T, Walzel A-V, and Wiermann B. 2020. NFDI4Culture - Consortium for research data on material and immaterial cultural heritage. Research Ideas and Outcomes 6: e57036. https://doi.org/10.3897/rio.6.e57036. See also: https://nfdi4culture.de/resources/user-stories.html (Accessed 7 May 2022)

[13] Bayerische Akademie der Wissenschaften. 2021. Corpus der barocken Deckenmalerei in Deutschland (CbDD). https://deckenmalerei.badw.de/ (Accessed 18 Jan 2022)

[14] See: https://iiif.io/news/2022/01/11/new-3d-tsg/ (Accessed 12 March 2022)

[15] Source: https://openrefine.org/ (Accessed 18 Jan 2022)

[16] See: https://gitlab.com/nfdi4culture/ta1-data-enrichment/openrefine-reconciliation-service-box (Accessed 7 May 2022)

[17] Rossenova, L. and Sohmen, L. 2021. OpenRefine to Wikibase: Data Upload Pipeline. *Wikiversity.* Available from: https://en.wikiversity.org/wiki/OpenRefine_to_Wikibase:_Data_Upload_Pipeline (Accessed 7 May 2022)

[18] Fauconnier, S., Espenschied, D., Moulds, L. and Rossenova, L. 2018. Many Faces of Wikibase: Rhizome's Archive of Born-Digital Art and Digital Preservation. *Wikimedia Blog. 2018.* https://wikimediafoundation.org/news/2018/09/06/rhizome-wikibase/. (Accessed 7 May 2022)

[19] Shorland, A. (2021). WBStack Infrastructure. Available from: https://addshore.com/2021/04/wbstack-infrastructure-2021/ (Accessed 12 March 2022)

[20] Alipio, S. and Shorland, A. (2021). Pre-launch Announcement and Preview of Wikibase.Cloud. *WikidataCon 2021, Online, 29-31 Oct 2021.* Available from: https://www.youtube.com/watch?v=GJ8rE-7F-zs (Accessed 12 March 2022)

[21] See: https://angularjs.org/ (Accessed 12 March 2022)