

# Integrating citizen experiences in cultural heritage archives: requirements, state of the art, and challenges

ENRICO DAGA\*, The Open University

LUIGI ASPRINO, University of Bologna

ROSSANA DAMIANO, University of Turin (CIRMA and Dipartimento di Informatica)

MARILENA DAQUINO, University of Bologna

BELEN DIAZ AGUDO, Complutense University of Madrid

ALDO GANGEMI, CNR, ISTC - Institute of Cognitive Sciences and Technologies

TSVI KUFLIK, The University of Haifa

ANTONIO LIETO, University of Turin (CIRMA and Dipartimento di Informatica) and ICAR-CNR

MARK MAGUIRE, Irish Museum of Modern Art

ANNA MARIA MARRAS, University of Turin

DELFINA MARTINEZ PANDIANI, University of Bologna

PAUL MULHOLLAND, The Open University

SILVIO PERONI, University of Bologna

SOFIA PESCARIN, CNR ISPC - Institute of Heritage Science

ALAN WECKER, The University of Haifa

Digital archives of memory institutions are typically concerned with the cataloguing of artefacts of artistic, historical, and cultural value. Recently, new forms of citizen participation in cultural heritage have emerged, producing a wealth of material spanning from visitors' experiential feedback on exhibitions and cultural artefacts, to digitally mediated interactions like the ones happening on social media platforms. Citizen curation is proposed in the context of the European project SPICE - Social

---

\*Corresponding author.

---

Authors' addresses: Enrico Daga, The Open University, Walton Hall, Milton Keynes, United Kingdom, MK6 6AA, enrico.daga@open.ac.uk; Luigi Asprino, University of Bologna, via Zamboni, 32, Bologna, Italy, 40126, luigi.asprino@unibo.it; Rossana Damiano, University of Turin (CIRMA and Dipartimento di Informatica), Corso Svizzera 185, Turin, Italy, 10149, rossana.damiano@unito.it; Marilena Daquino, University of Bologna, via Zamboni, Bologna, Italy, 40126, marilena.daquino2@unibo.it; Belen Diaz Agudo, Complutense University of Madrid, Facultad de Informática. Ciudad Universitaria, Madrid, Spain, 28040, belend@ucm.es; Aldo Gangemi, CNR, ISTC - Institute of Cognitive Sciences and Technologies, Rome, Italy, aldo.gangemi@cnr.it; Tsvi Kuflik, The University of Haifa, Mount Carmel, Haifa, Israel, 31905, tsviak@is.haifa.ac.il; Antonio Lieto, University of Turin (CIRMA and Dipartimento di Informatica) and ICAR-CNR, Corso Svizzera 185, Turin, Italy, 10149, antonio.lieto@unito.it; Mark Maguire, Irish Museum of Modern Art, Dublin, Ireland, mark.maguire@imma.ie; Anna Maria Marras, University of Turin, Turin, Italy, annamaria.marras@unito.it; Delfina Martinez Pandiani, University of Bologna, Department of Computer Science and Engineering, Bologna, Italy, 40126, delfinasol.martinez2@unibo.it; Paul Mulholland, The Open University, Walton Hall, Milton Keynes, United Kingdom, MK7 6AA, paul.mulholland@open.ac.uk; Silvio Peroni, University of Bologna, via Zamboni, Bologna, Italy, 40126, silvio.peroni@unibo.it; Sofia Pescarin, CNR, ISPC - Institute of Heritage Science, Florence, Italy, sofia.pescarin@cnr.it; Alan Wecker, The University of Haifa, Mount Carmel, Haifa, Israel, 31905.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, or post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

XXXX-XXXX/2021 (Pre-print)/1-ART0 \$15.00

<https://doi.org/10.1145/3477599>

Participation, Cohesion, and Inclusion through Cultural Engagement - as a methodology for producing, collecting, interpreting, and archiving people's responses to cultural objects, with the aim of favouring the emergence of multiple, sometimes conflicting viewpoints, and motivating users and memory institutions to reflect upon them. We argue that citizen curation urges to rethink the nature of computational infrastructures supporting data management of memory institutions, bringing novel challenges that include issues of distribution, authoritativeness, interdependence, privacy, and rights management. To approach these issues, we survey relevant literature towards a distributed, Linked Data infrastructure, with a focus on identifying the roles and requirements involved in such an infrastructure. We show how existing research can contribute significantly in facing the challenges raised by citizen curation, and discuss challenges and opportunities from the socio-technical standpoint.

CCS Concepts: • **Information systems** → **Digital libraries and archives**; Collaborative and social computing systems and tools; • **Applied computing** → **Arts and humanities**.

Additional Key Words and Phrases: cultural heritage, semantic web, citizen curation

## 1 INTRODUCTION

Digital archives of memory institutions are typically concerned with the cataloguing of artefacts of artistic, historical, and cultural value. Nonetheless, in recent years new forms of citizen participation in cultural heritage have emerged, producing a wealth of material relevant to curatorial practices, spanning from visitors' experiential feedback to exhibitions and cultural artefacts, to digitally mediated forms of interactions like the ones happening on social media. As a consequence, how to integrate citizens' contributions in curatorial practices has become a growing subject of interest in several research fields.

Citizen curation is proposed in the context of the European project SPICE - Social Participation, Cohesion, and Inclusion through Cultural Engagement - as a methodology for producing, collecting, interpreting, and archiving people's responses to cultural objects, with the aim of favouring the emergence of multiple, sometimes conflicting viewpoints, and motivating the users and memory institutions to reflect upon them. One of the main assumptions of SPICE is that, although social media platforms allow heritage institutions to connect with the public, they present significant limits when used for participatory cultural activities. Traditionally, the idea of 'participatory culture' is associated with initiatives that gather and analyse users' engagement with cultural heritage through social media platforms [40], such as Instagram, Twitter, and Facebook [57] [2]. However, heritage institutions may wish to share digital objects online, to track their provenance and reuse, access to social media history, and so on. The heterogeneity of platforms, API services, and terms and conditions, hampers the creation of a social space where museum artefacts and resources coexist with the responses of visitors in meaningful way. Particularly, social media raise concerns in terms of privacy and usage tracking. First, content published on social media is stored in an infrastructure that is not controlled by the cultural institution. Second, social media are currently used by cultural institutions as a broadcasting media in their communication and marketing strategy, rather than as a means to develop one-to-one interpretive experiences between an artwork and an individual. We argue that there is a fundamental imbalance in the power relation between content producers and social media service providers, where the latter make all efforts to enable free user expression but reject liability for the messages (interpretations) that are generated through the platforms [41].

Equally, we argue that digital archives of memory institutions are not prepared for integrating the experience of citizens. The recent effort of the EU in supporting an European Research Infrastructure for Heritage Science, clearly shows the interest for developing a shared infrastructure that supports researchers and cultural heritage institutions [6]. Since the CHARISMA project from 2009 to 2014 (Cultural heritage Advanced Research Infrastructures: Synergy for a Multidisciplinary Approach to Conservation/Restoration), this effort has continued with DARIAH (Digital Research Infrastructure for the Arts and Humanities [www.dariah.eu](http://www.dariah.eu)), CLARIN and CLARIN PLUS 2008 - 2017 (Common Language Resources and Technology Infrastructure structure, [www.clarin.eu](http://www.clarin.eu)), CENDARI 2012-2016 (Collaborative European Digital Archive Infrastructure, [www.cendari.eu](http://www.cendari.eu)), and now E-RIHS (The European Research Infrastructure for Heritage Science) and IPERION HS 2020-2023 (Integrating Platforms for

the European Research Infrastructure ON Heritage Science). In case of E-RIHS, the goal is to deliver integrated access to expertise, data and technologies through a standardised approach (<http://www.e-rihs.eu>), implementing Open Access EU strategy<sup>1</sup>, participating in EOSC and contributing to the FAIRification of heritage data<sup>2</sup>. More importantly, an infrastructure should be able to address requests of emerging research, such as in SPICE, so as to support big data storage, analysis, knowledge extraction and reuse, and meet requirements of experts (researchers and curators), businesses in cultural industry, and the broader public (for instance by promoting reflection, knowledge acquisition, engagement).

Citizen curation requires a technical infrastructure that is *egalitarian in essence*, where memory institutions, citizens, scholars, and businesses share the “power” over the assets they produce and the associated metadata. But what does this mean from the technological standpoint? And how to frame it in the context of Web technologies? In the light of this research program, it is an open question what type of technologies and systems could support management and preservation of data produced by citizen curation. To answer this question, in this article (i) we characterise citizen curation from the point of view of user roles and (ii) we devise a general workflow. From the resulting scenario, (iii) we derive a set of socio-technical requirements that an infrastructure for citizen curation should satisfy. We conclude that a technical solution for citizen curation will necessarily be a composition of multiple technologies in a distributed ecosystem.

The main contribution of this article can be summarised in the following research questions: *What are the requirements of citizen curation? What are the technologies that contribute to an infrastructure for citizen curation? To what extent existing systems support cooperation? Are there any missing critical components in the state of the art? What types of connections need to be established? What approach should we take in order to fill the gaps?* To address these questions, we survey relevant research in computational ecosystems for cultural heritage focusing on technologies and tools that could contribute to build social applications for citizen curation, and tackle issues of distribution, authoritativeness, interdependence, privacy, and rights management.

In Section 2 we introduce *citizen curation*. From this notion, we characterise our *problem space*, by defining a general workflow, user roles, and the requirements for a computational infrastructure for citizen curation (Section 3). In Section 4 we look at the *solution space*, by surveying technologies grouped in research areas that match with aforementioned requirements. In the light of the survey, in Section 5 we discuss the state of the art, challenges and opportunities. In Section 6 we conclude the article.

## 2 BACKGROUND: CITIZEN CURATION

Traditionally, museums could be thought of as providing an authoritative account of their collection, informing and educating citizens as to the meaning, importance and relevance of their artefacts. The role of the citizen was to appreciate the artefacts and acquire the knowledge and stories associated with them. The role and purpose of museums now tends to be viewed quite differently. The Faro convention on the value of cultural heritage for society [31] declares the need to “involve everyone in society in the ongoing process of defining and managing cultural heritage”, that “every person has a right to engage with the cultural heritage of their choice”, and that “all cultural heritages [should be treated] equitably and so promote dialogue among cultures and religions”. The International Council of Museums (ICOM)<sup>3</sup> defines [56] the museum as an institution which “acquires, conserves, researches, communicates and exhibits the tangible and intangible heritage of humanity and its environment for the purposes of education, study and enjoyment”. This definition can be perceived as consistent with a more traditional view of the museum as having a responsibility to communicate an understanding of heritage to the public. A proposed revision in 2019 to the ICOM definition [56] describes museums as “democratising, inclusive

<sup>1</sup><http://ec.europa.eu/research/openscience/index.cfm?pg=openaccess>

<sup>2</sup><https://op.europa.eu/en/web/eu-law-and-publications/publication-detail/-/publication/5253a1af-ee10-11e8-b690-01aa75ed71a1>

<sup>3</sup><https://icom.museum/en/>, accessed 13/05/2021.

and polyphonic spaces for critical dialogue about the pasts and the futures". It goes on to state that museums are "participatory and transparent, and work in active partnership with and for diverse communities to collect, preserve, research, interpret, exhibit, and enhance understandings of the world". This proposal stimulated a real debate inside the museum community for this reason it was decided to work on a new proposal, through a new participatory process, which will be voted on during the ICOM General Assembly which will take place in Prague in May 2022. The Faro convention and the proposal to revision to the ICOM museum definition both extend the traditional conceptualisation of the museum in two ways. First, they both highlight that there is not necessarily a single interpretation of heritage. There may be multiple interpretations. Second, they emphasise that the role of the citizen is not confined to acquiring what is presented to them. Citizens can be actively engaged in sharing their voices, participating in dialogue and creating understandings. This trend toward multiple voices and active participation can be seen in recent initiatives to decolonise the museum, challenge the dominant narrative and introduce new perspectives.

Within the SPICE project we are developing tools and methods to support a process we term Citizen Curation. We define Citizen Curation as *citizens applying curatorial methods to archival materials available in heritage and memory institutions as well as to items depicted in exhibitions in order to develop their own interpretations, share their own perspective and appreciate the perspectives of others*. Crucially, our definition of Citizen Curation covers both citizens sharing their own perspectives and also engaging positively with the interpretations of others. The aim is not to just provide multiple interpretations so that the citizen can select the one that fits with their World view, but rather to promote dialogue across perspectives as anticipated by the Faro convention.

Our definition of Citizen Curation has deliberate parallels to the concept of empathy. Zaki [119], characterises empathy as encompassing a number of ways in which people respond to each other: identifying what the other person feels (cognitive empathy), sharing the emotion of the other person (emotional empathy) and wanting to improve the experiences of the other person (empathic concern). Empathy can often come easily toward people similar to oneself; people are tribal by nature [5]. However, empathy can also be cultivated toward perceived out-groups. Two processes that can help to build empathy toward out-groups are perspective giving (sharing one's point of view) and perspective taking (seeing the World from someone else's perspective) [5]. Citizen Curation, in incorporating both the sharing and appreciation of perspectives, recognises the role that museums can potentially play in building empathy and cohesion across as well as within communities.

Our definition is informed by previous initiatives that have engaged citizens in the curatorial process. Mauer [72] and Hill et al. [53] characterise Citizen Curation as a process in which citizens with little or no background in museum curation are provided with training and guidance to create their own physical and virtual exhibitions. Our approach builds on this work but aims to extend Citizen Curation to larger scale participation without additional training. Some previous work has used online tools to widen participation in Citizen Curation. Moqtaderi [75], uses the term citizen curator to describe members of the public voting for an artwork to be included in an exhibition curated by the museum. The citizen curator initiative developed by Ride [90], involved citizens sharing contributions via Twitter, later used in a video installation developed by the museum. We take a related approach but emphasise the importance of engaging citizens in perspective taking (appreciating the viewpoints of others) as well as perspective giving. The case studies within the SPICE project develop and apply Citizen Curation in a range of contexts. For example, the case study at the Irish Museum of Modern Art (IMMA) will support groups less able to visit the museum physically, for example due to illness, to share their perspectives online. Citizens will be encouraged to think about universal personal themes such as family to make interconnections across groups. Hecht Museum in Israel will enable members of religious and secular communities, in particular minority communities, to express and share their viewpoints in order to find commonalities and also understand differences. Other case studies in the SPICE project work with older people, people living in rural communities, Deaf people, and children. A common thread running through the studies is the use of Citizen Curation to both share and appreciate perspectives in order to build understanding and empathy.

### 3 REQUIREMENTS ANALYSIS: THE PROBLEM SPACE

Citizen curation covers an extensive variety of issues from different points of view, including the ones of knowledge representation and human-computer interaction. Here, we focus on requirements related to data management with a special attention to aspects related to distribution. The following requirements have been designed during the first six months of the SPICE project in a number of discussion groups involving technologists, designers, and professionals in the cultural heritage sector. We will see how many of these requirements are related to the handling of rights, access control, and monitoring of use, of museums' digital assets, which emerged to be critical in relation to the development of citizen curation scenarios.

#### 3.1 Integrating citizen experiences in cultural heritage archives

The following guide scenario illustrates a typical citizen curation pipeline that could be implemented with the aid of systems such as digital archives, Web sites, and social media:

*Cath is a museum curator. She decides to run an online activity which supports citizens in sharing and reading personal stories inspired by artworks. She selects a set of artworks to be used in the activity. This involves checking that the museum has appropriate permissions to use images of the artworks in the activity and, where necessary, securing appropriate permission from the rights holder. Once the activity has been prepared it is launched on a website developed by a company of the cultural industry sector. Citizens can choose to take part in the activity anonymously, create an account on the system or login via a third-party, mobile application. Citizens can select one of the artworks, tell a personal story related to the artwork and send it to a friend. The friend can send a response to the person who wrote the story. The citizen can also choose to share the story with the curator. Even when a citizen has decided to share their story with the curator, they retain the option to withdraw the story at any time. Cath is able to monitor stories contributed by the citizens. Any story that may contain inappropriate content is automatically flagged and she can choose to remove it. Once the activity had been running for two weeks, Cath creates an online exhibition featuring a selection of the contributions shared with her. In the presentation, she draws attention to the different ways in which artworks have been interpreted. She decides to close the activity to further contributions and relaunch it with a new set of artworks later in the year. Finally, she curates the contributed content and include it in the museums digital archive for preservation.*

The above scenario reflects a number of needs of the museum sector that are reflected in the literature. Many museums wish to actively engage citizens and encourage them to share their own contributions with the museums. Illustrative prior examples of citizen participation are discussed in section 2 (e.g. [72], [53], [75] and [90]). Museum participation initiatives not only aim to support citizens in sharing their contributions with the museum but also each other. For example, Spence et al [102] describe an app that can be used to share museum-inspired stories with friends. When museums wish to engage with citizen online, they often turn to social media platforms due to their availability. However, these raise additional challenges. For example, Ride [90] in using social media for citizen contributions, had to navigate Intellectual Property issues related to the rights associated with shared content and rights given to the platform. The use of social media by museums in practice, is often one-way as a broadcasting medium. Lazzeretti et al [65] propose that this lack of interaction may be due the museum potentially being seen as accepting responsibility for, or condoning, opinions found in user generated content which it may be difficult for the museum to track and moderate. Therefore, the above scenario serves to capture a range of previously documented challenges related to museum engagement with citizens.

From the above scenario, we develop an abstract workflow for citizen curation, that we then use as a reference for describing user level requirements. If we look into the socio-technical context where digital systems for citizen engagement with cultural heritage such the one envisaged in our scenario are developed, it becomes evident how the traditional distinction between data provider and data consumer is insufficient. Instead, we can identify four major roles: the owner, the custodian, the builder, and the end user.

**The owner** is the copyright holder of the cultural heritage asset that is selected for the citizen engagement activity. Owner can be the public body having ownership of the archaeological site (e.g. a National Ministry of Culture) or the artist that created the sculpture or her heirs. In practice, this role can be played by a representative organisation, such as the rights management society delegated to act on behalf of the artist. Main concern of this role are the control on the way the asset is exploited by the community. For example, the owner may want some conditions to be applied on the way an artwork is displayed. Interestingly, this role is not different from the one of users of Web social media and their concerns with relation to privacy (we will expand on this connection). However, the owner is rarely involved in the management of the asset, that is typically delegated to another entity, that we call here *custodian*.

**The custodian** is the organisation delegated to preserve and valorise the asset. Cultural heritage archives are maintained by custodians: museums, national archives, and so on. Concerns of the custodian are the authenticity of the asset, the quality of the metadata associated, and the management of a number of transactions with third-parties interested in accessing and using the asset. Although the valorisation of the asset is part of the mission of custodians, the development of media technologies for supporting access and engagement are typically delegated to another type of entity, for example, a company that *builds* systems for the cultural industry sector.

**The builder** is the entity developing means for supporting engagement activities with cultural heritage, for example an interactive tour guide deployed as a mobile app, a virtual exhibition accessible on the Web, or a tool used in a participatory design workshop organised with a local community. These systems create the conditions for engagement and at the same time manage the identities of the users involved, for example, by recording the personal email accounts to the content generated with a visitor's mobile app, or by linking to other social media accounts. This model reflects well the type of activities that museums perform on social media. When social media are used for cultural campaigns, a *custodian* is delegating to a *builder* (e.g. Twitter) the mediation with the citizens. The process described so far is illustrated in Figure 1. However, all of the above issues are not exclusive




Fig. 1. Citizen curation: from artists to citizens.

of the relationship between artists, museums, media companies, and citizens.

**The end user** is now engaging with cultural heritage proactively, for example through mass social media, producing a wealth of content that is increasingly gaining the attention of memory institutions. Interestingly, the aforementioned issues also apply to user-generated content, when such content is treated as a first-class citizen of cultural heritage archives. In typical social media platforms, the content produced by users is often legally owned by the company who runs the service. Although when using YouTube, Facebook, Instagram, and Twitter the users are technically in control of the content that they upload, they are not in control of how their

content is used. The companies running the service typically request and receive comprehensive licences to reuse, distribute, and profit from the content we share on the system. As a result, the companies may be free to reuse photos of our holidays as a context for their advertisements, purpose your comments in a different context, or even sell it to third parties. How can users be empowered and choose how their content is used? How can users express how they want to be considered with respect to the copyright law that applies to the produced content?

The model introduced in Figure 1 can be easily reversed, as in Figure 2. For instance, a visitor shoots a selfie next to an artwork and shares it on Instagram, tagging the museum. She is the *owner* of the asset, and the company controlling the social media becomes the *custodian*. Moreover, memory institutions today want to interact with such wealth of material, connect it to the cultural heritage archive, and make it available in the context of cultural and social studies. In this alternative scenario, the end user is the archivist or the scholar that explores the museum archive and *makes sense* of the relationships between the archival content and citizen experiences. This reverted scenario is illustrated in Figure 2.

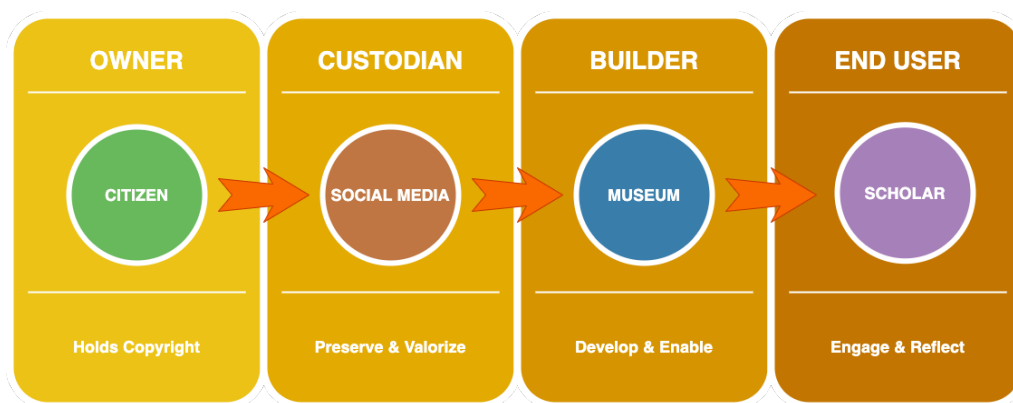


Fig. 2. Citizen curation: from citizens to scholars

### 3.2 Requirements of an ecosystem for a citizen curation

In the following paragraphs we describe requirements defined with respect to the guiding scenario introduced in Section 3.1 and the four roles identified in the model. We organise functional requirements in relation to the roles, then introduce non-functional requirements by focusing on the interaction between them.

**The owner.** From the point of view of the owner - whether an artist being the copyright holder on an image of my artwork, or a citizen whose data or opinions have been collected by a company - the main concern is having agency with respect to the terms of access and use of the asset. This includes describing how the asset can be used and granting or revoking rights to target users or organisations:

- R 1. owner: express *Ability to express the terms and conditions of use of an asset*
- R 2. owner: rights *Ability to grant or revoke licence or rights associated with a digital asset*

A rights management organisation is often delegated by the owner for negotiating terms of use with the museum and possibly third parties involved. In addition, a owner could ask the museum for a report on how assets were used, for example, in a given public event:

- R 3. owner: negotiate *Ability to participate in the negotiation of digital rights with third-parties*

R 4. owner:know *Ability to know how the digital asset is used*

R 5. owner:claim *Ability to request or claim ownership of an asset*

**The custodian.** The entity delegated of managing cultural heritage assets is the *custodian*. A custodian wants to be able to interact with the owner in order to define and manage permissions on assets. To this extent, requirements 1-3 apply also to the custodian. Secondly, citizen curation requires the curator to handle a data management life cycle, including the set up of repositories, the curation of metadata, the organisation of those into collections, the management of access control on assets, and the monitoring of data usage:

R 6. custodian:manage *Manage a repository of assets and associated metadata*

R 7. custodian:curate *Curate an open-ended metadata set, according to the organisation mission and needs, organise and present the content into collections, and guarantee the quality of information*

R 8. custodian:publish *Publish a repository of assets and associated metadata to third-party applications*

Crucially, curators want to be able to express and manage permissions on data to be granted to third parties:

R 9. custodian:visibility *Manage the visibility of repositories by third parties, including publishing metadata as open data*

A set of requirements derive from the role of museums as mediators of the copyright holders in managing the access and usage of the digital asset:

R 10. custodian:copyright *Declare the intellectual property (copyright) associated with the assets*

R 11. custodian:enable *Enable third-parties to request access to the assets and/or its metadata*

R 12. custodian:access *Manage access control to the digital assets and metadata, including the ability to grant or revoke access to a digital asset*

Since memory institutions want to enrich their digital archives with traces of citizen experiences, user-generated assets become first class citizens of the digital heritage. This implies adding new metadata on opinions about the artworks, including conflicting viewpoints, alongside authoritative metadata:

R 13. custodian:viewpoints *Integrate in the archive different viewpoints alongside authoritative metadata*

As a result, custodians must cope with privacy issues. Identify and filter citizens' sensitive or inappropriate contents is a primary concern.

R 14. custodian:monitor *Monitor content integrated into the archive to raise issues with relation to copyright infringement or privacy law*

R 15. custodian:inappropriate *Monitor content integrated into the archive to raise issues with relation to sensitive and inappropriate content*

However, digital assets that are subject of citizen curation may come from multiple sources. Therefore, it is of crucial importance that digital archives can be linked to external repositories:

R 16. custodian:link *Link the metadata to content provided by other stakeholders*

Finally, museum organisations want to collect information on the usage of digital assets, therefore supporting scholars in studying how cultural heritage is perceived and copyright owners in receiving feedback:

R 17. custodian:usage *Monitor, trace, and analyse the usage of the assets by third parties*



In perspective, recording usage traces will allow to reason over the data journeys [3, 26] associated with cultural assets and support use cases such as policy propagation: where a condition or term of use should be transferred to derived assets [25]. Besides, the history of access and use of digital assets may be itself an object worth to be considered by memory archives facing the challenge of dealing with born-digital cultural heritage [70]. The distributed online infrastructure should have means for recording the provenance of digital assets in a way that activities can be recorded for future inspection by allowed parties [22].

**The builder.** A builder - for instance, a company that wants to access cataloguing data and social media data to develop a tour guide application - values technologies that are reusable, reliable (in terms of data security and storage), and interoperable with the repository of custodians. Builders need to rely on an infrastructure that allows them to request permissions to accessing the data, obtaining means for secure access (e.g. an API key), and perform operations such as reading data. These requirements hold in both of the above introduced models (i.e., the standard and the *reversed* one, when the custodian is the social media provider):

R 18. builder:interop *Inter-operate with data managed by the custodian*

R 19. builder:request *Request permissions to use digital assets controlled by the custodian, and obtain access*

The receiver of the rights should be supported in answering questions such as *why a given resource is not accessible?* or *what can be done to obtain access?* This includes explaining conditions, permissions, and prohibitions (e.g. *why the system was able to download the digital image in the past but cannot do it today?*):

R 20. builder:viewterms *Review the terms of use associated to a digital assets*

However, digital objects propagate on the Web very easily as computing machines operate essentially by copying and exchanging data objects. Therefore, it may not be possible to technically avoid misuse. Instead, we can ensure that rights are expressed in a way that potential breaches can be described and explained:

R 21. builder:explain *Receive an explanation on why a digital resource can(not) be accessed or used for a given purpose*

Finally, a builder benefits from an ecosystem of data sources and services that, while being provided by a multiplicity of vendors, should rely on open standards:

R 22. builder:standards *Design solutions relying on open standards rather than vendor-specific technologies*

By providing a service to an audience, the builder mediates between the digital archives and the target users.

**The end user: citizens and scholars.** End users engage with cultural heritage content and produce new interpretations (e.g. a museum visitor engages with the artwork using an augmented reality app, or a scholar studies citizens' contribution archived by the museums). In this work we don't focus on how user tasks are performed, that is, interaction design aspects. Rather, we select a set of high-level tasks that users face when interacting with the content of the archive, such as *exploration*:

R 23. enduser:explore *Ability to explore the contents of the digital archive*

Citizen curation is performed by *sharing* contents relevant to their opinions, *confronting* artefacts, and user-generated content (e.g. opinions, stories, tags, images), and *enhancing* these with personal opinions:

R 24. enduser:share *Signpost the digital material to other users and share with them opinions and viewpoints*

R 25. enduser:compare *Critically compare and evaluate digital material originating from multiple sources*

R 26. enduser:enhance *Annotate and enrich the digital assets with other content (as personal interpretations)*

User-generated content may include information of sensitive nature, including personal identifiable information. An infrastructure supporting the exchange of data of such nature must be equipped with methods for notifying users on these issues:

R 27. `enduser:privacy` *Receive support in knowing whether the shared content could be a threat to privacy*

Similarly, users should be aware of the copyright information related to the digital assets:

R 28. `enduser:copyright` *Understand information about copyright and terms and conditions associated to the digital assets in an accessible way*

Finally, citizens should be recognised as the authors of their contributions:

R 29. `enduser:claim` *Declare ownership and claim intellectual property on the contributed content*

By declaring the ownership and property of the produced content, citizens become *owners*, and the loop of citizen curation is closed.

**The ecosystem.** In addition to the role-based requirements introduced so far, we identified a set of *non-functional* requisites that an ecosystem for citizen curation needs to support. In a citizen curation scenario, third-party applications interact with content published by digital archives of memory institutions and user-generated contents, which are shared with multiple organisations and users. Personal identities do not need to be shared across multiple contributions - if the user does not want to. The infrastructure should be able to support partnerships and cross-museums initiatives, enabling them to share both digital content and citizens' contributions. Primarily the ecosystem requires means for linking and consuming data from distributed sources:

R 30. `eco:link` *The infrastructure should enable digital archives to link data to external sources*

R 31. `eco:consume` *The infrastructure should enable systems to consume data from external digital archives*

R 32. `eco:collect` *The infrastructure should support digital archives in collecting data from third-party applications, for example, receiving notifications when content is created or updated*

Secondly, the ecosystem should be distributed, avoiding a central authority managing agents' and users' identities. Organisations and users should be allowed to delegate third-parties in hosting authoritative identifiers for the entity mentioned in the metadata:

R 33. `eco:identity` *The infrastructure should have means for distributed identity management*

R 34. `eco:entity` *The infrastructure should have means for delegating data identities to third-party authorities*

Several role-based requirements discussed imply that organisations and users delegate other entities to act on their behalf, for example, for managing owned resources and related access control:

R 35. `eco:agency` *The infrastructure should have means for supporting the delegation of agency*

Similarly, digital archives may be dynamically populated by third-party systems:

R 36. `eco:manage` *The infrastructure should have means for supporting data management operations by third-party applications*

Lastly, the ecosystem of tools should be reusable and open so as to enable the development of applications:

R 37. `eco:extend` *The infrastructure components should be extensible and reusable with an open source licence*

## 4 RELATED WORK: THE SOLUTION SPACE

In this section we survey technologies that potentially contribute to the development of an ecosystem for citizen curation. Our survey methodology is based on the following actions. From the role-based analysis of requirements, we identified six areas of interest, namely: (a) Data management for cultural heritage; (b) Web technologies for metadata publishing and exchange; (c) Rights data management; (d) Crowd-sourcing and cultural heritage; (e) Social Media uses in cultural heritage; and (f) Distributed online social networks. The mapping between requirements and areas is summarised in Table 1. Secondly, we identify one or more recent surveys of technologies relevant to each area and we discuss each technology in the light of the requirements. We focus on three type of technologies for data management: (a) end-user tools, (b) services, and (c) components meant to be used within a distributed technical infrastructures. When surveyed technologies address significantly different requirements we describe those separately so as to address their peculiarities. For the sake of readability, in a few cases (see Section 4.5 and 4.6) we discuss technologies all together, as these address the same requirements. We do not survey related work on approaches to co-design, engagement, metadata management, vocabularies, datasets, or end-user applications that are not targeted to data acquisition and management, that do not have an element of distribution, or that are not maintained anymore. While we do not claim that the survey is exhaustive, we expect it to significantly represents current state-of-the-art solutions for the management of digital heritage.

### 4.1 Data management

In the last years, collection management platforms have become an essential part of the dissemination plans of cultural institutions. Situated midway between Content Management Systems and professional cataloguing tools, these hybrid systems combine the traditional functions of museum software with the need to make collections available online. Solutions range from proprietary high-end products such as the TMS Suite by Gallery Systems<sup>4</sup> to open source platforms. Based on previous surveys and comparisons [49, 117], in the following we review Omeka S, DSpace and Fedora (with its spin-offs Islandora and Samvera).

**Omeka S**<sup>5</sup> has established itself as a solution for small to medium projects, thanks to its affordable requirements and ease of use [71]; its modular architecture facilitates the integration with external services such as the ones implementing the International Image Interoperability Framework (IIIF) and indexing services. While Omeka Classic relies on the DCMES, its linked data version, Omeka S, leaves the user free to create her/his own metadata schema from any RDF vocabulary [67], allowing the same set of semantically described items to be shared among different sites on the same installation and across different installations through APIs. From this perspective, Omeka partly satisfies the requirements concerning interoperability and openness for the role of the builder (such as requirements R18/builder:interop and R22/builder:standards). Omeka's modular structure lends itself to data source linking (requirements R30/eco:link, R31/eco:consume, but it lacks predefined policies and protocols to regulate them as prescribed by R32/eco:collect or R33/eco:identity apart from those embedded in specific connection modules (e.g., modifications to items in an external Fedora repository would determine updates in the items imported into Omeka S). Omeka S can support the acquisition of user generated content as part of its workflows, through dedicated modules. Requirements R6/custodian:manage to R12/custodian:access could be easily met. By converse, a core requirement for citizen curation such as R13/custodian:viewpoints could not be straightforwardly enforced in Omeka S. From the enduser's perspective, Omeka S enables the curator to create complex views (R25/enduser:compare), but the interactive creation and enrichment of personal views is not enabled.

**DSpace** [100] is an open-source platform for collecting, managing, indexing, and distributing digital assets including text, images, videos and data sets. DSpace allows to search and retrieve items, to upload digital

<sup>4</sup><https://www.gallerysystems.com/solutions/collections-management/>

<sup>5</sup><https://omeka.org/s/>

Table 1. Requirements and related research areas

R	Data M.	Web T.	Rights M.	Crowd-S.	Social Media	DOSN
R1/owner:express			✓			
R2/owner:rights		✓	✓			
R3/owner:negotiate			✓			
R4/owner:know			✓			
R5/owner:claim			✓	✓		
R6/custodian:manage	✓					
R7/custodian:curate	✓					
R8/custodian:publish	✓	✓				
R9/custodian:visibility	✓	✓	✓			✓
R10/custodian:copyright	✓		✓			
R11/custodian:enable	✓	✓	✓			
R12/custodian:access	✓	✓	✓			✓
R13/custodian:viewpoints	✓			✓		
R14/custodian:monitor	✓	✓				✓
R15/custodian:inappropriate	✓	✓				
R16/custodian:link		✓				
R17/custodian:usage		✓				
R18/builder:interop	✓					
R19/builder:request		✓	✓			
R20/builder:viewterms			✓			
R21/builder:explain			✓			
R22/builder:standards	✓	✓	✓	✓	✓	✓
R23/enduser:explore				✓		
R24/enduser:share					✓	✓
R25/enduser:compare				✓	✓	✓
R26/enduser:enhance				✓		
R27/enduser:privacy	✓	✓	✓			✓
R28/enduser:copyright	✓		✓			
R29/enduser:claim			✓	✓	✓	
R30/eco:link	✓	✓				
R31/eco:consume	✓	✓				
R32/eco:collect	✓	✓				
R33/eco:identity		✓				✓
R34/eco:entity		✓				
R35/eco:agency		✓				
R36/eco:manage		✓				
R37/eco:extend	✓	✓	✓	✓	✓	✓

items and to decide user roles. It complies with many protocols for access, ingest, and export (such as: OAI-PMH, WebDAV, RSS, ATOM) and supports common authentication methods, including: LDAP (and hierarchical LDAP), Shibboleth, X.509, IP-based. DSpace covers various requirements for custodians (i.e. requirements R6/custodian:manage, R7/custodian:curate, R8/custodian:publish, R9/custodian:visibility,

R10/custodian:copyright), builders (i.e requirements R18/builder:interop and R22/builder:standards), for end users (i.e. R23/enduser:explore and R28/enduser:copyright) and for the whole ecosystem (i.e. requirement R37/eco:extend).

**Fedora** [83] is an open-source repository for long-term digital preservation and reliable access of any type of digital objects. Fedora is designed to fulfill the requirements of interoperability and extensibility. As far as interoperability concerned, Fedora is standards-based since it supports a number of widely adopted standards. Specifically, it provides a robust RESTful API layer, it allows to serve data as RDF, thus meeting many requirements from the ecosystem's and custodian's perspective (R6/custodian:manage, R7/custodian:curate, R8/custodian:publish, R9/custodian:visibility, R10/custodian:copyright, R31/eco:consume, R32/eco:collect). Fedora is designed to integrate with other applications and services to provide additional functionalities (such as dissemination using OAI-PMH, deposit with SWORD etc.) thus covering the requirement R37/eco:extend. Fedora can control access to content via a pluggable framework compliant with SOLID specification, according to requirement R33/eco:identity and R34/eco:entity. Fedora implements the memento protocol for enabling the versioning of the resources. Finally, concerning security issues, Fedora is able to integrate with existing authentication systems such as LDAP and Shibboleth.

**Islandora and Samvera** The open source Islandora and Samvera projects share the use of Fedora as repository system, to which they add web publication and user management functions. So, they both inherit the advantages and limitations of Fedora, while their support to the requirements not met by Fedora depend on the specific additional components of each system (and, in the case of Samvera, of the specific configuration of installed modules). Islandora<sup>6</sup> software framework plugs a well-known CMS, Drupal, into Fedora, thus augmenting it with content publishing and user management functions. Open to different content types, from scientific to cultural data, Islandora is characterised by a focus on access through advanced search functions, thanks to the integration with Solr; similarly to Omeka S, it can be integrated with external servers such as IIF and OAI-PMH servers. The integration of Fedora with content management functions extends the framework's capability to comply with citizen curation processes: in particular, it enables the creation of collections on top of the items in the catalogue, thus partly addressing the requirement R13/custodian:viewpoints; thanks to the user management system, it allows tailoring the access and permissions over the items and collections to specific role types (R12/custodian:access. For the end user, the availability of search tools results in a more flexible, personal exploration of the repository, thus meeting the requirements concerning R23/enduser:explore, R25/enduser:compare; acquisition of user content can be accomplished through Drupal Form API, in accordance with the requirement R26/enduser:enhance. In the same spirit, Samvera<sup>7</sup>, an evolution of the former Hydra system, is an open source suite of repository software tools aimed at creating flexible solutions for specific knowledge domains and tasks. Also built on top of Fedora, this framework does not offer canned solutions for citizen curation, but sets of specific modules that can be in principle configured and programmed to satisfy the citizen curation requirements (thus meeting the requirements R37/eco:extend and R18/builder:interop). This openness, however, although paving the way to a creating a citizen curation system from this framework, is an obstacle to the easy development of installations, especially by smaller institutions. The collection management function in Samvera is provided by the Blacklight<sup>8</sup> search and discovery platform, which allows it to meet most end user requirements (R23/enduser:explore R25/enduser:compare, R26/enduser:enhance) and the custodian's need to create and publish collections (R7/custodian:curate,R8/custodian:publish) and manage their visibility (R9/custodian:visibility). Similarly to Islandora, Samvera's identity management system

<sup>6</sup><https://islandora.ca/>

<sup>7</sup><https://samvera.org/>

<sup>8</sup><http://projectblacklight.org/>

supports the restriction of access to the repository items (R12/custodian:access, but without providing a proper DRM management system).

**ResearchSpace** [77] is an open source platform for enabling researchers to create, link, share and search data by using Semantic Web languages and technologies. It provides an assertion and argumentation model for tracing multiple perspectives on historical facts; enables a multilevel visual representation of resources; allows the creation of data and narratives in form of knowledge graphs; captures provenance of data; expresses researchers' views as graphs connecting narratives, data, processes, and arguments. ResearchSpace builds on the knowledge graph platform, enabling customisation and extensibility of the interaction with the graph database through the use of Semantic Web standards and expressive ontologies for schema modelling based on CIDOC-CRM. The platform also integrates external tools including OntoDia, MIRADOR Image Viewer with an IIF Image Server. Due to its ability of capturing multiple viewpoints with Semantic Web standards, ResearchSpace addresses the requirements R13/custodian:viewpoints, R18/builder:interop and R22/builder:standards. Moreover, it also allows custodians to manage and curate metadata (cf. requirements R6/custodian:manage, R7/custodian:curate, R8/custodian:publish, R9/custodian:visibility, R10/custodian:copyright, and R16/custodian:link) and it enables end users to explore digital assets (R23/enduser:explore) and to read copyright terms (R28/enduser:copyright). The ecosystem requirements supported by ResearchSpace are the following: R30/eco:link, R31/eco:consume, R37/eco:extend.

A leading platform that contributes to making available digital cultural content to citizens is **Google Arts and Culture (GAC)**. It is especially focused on virtual exhibitions and digital collections, and to date it allows the exploration of the collections of more than 2000 cultural institutions from around the world [69]. GAC covers various requirements of citizens as end users, including R23/enduser:explore, R24/enduser:share, R25/enduser:compare, and R26/enduser:enhance. Citizens can explore millions of images of artworks, and compare them in multiple ways, such as confronting them by colour or historical period; create personal virtual galleries by choosing artworks and commenting on them; and share artworks or galleries with others. GAC does not, however, allow users to claim ownership over their personal galleries and interpretations (requirement R29/enduser:claim). Additionally, according to Lee et al. [66], GAC does not cover requirement R28/enduser:copyright thoroughly, as GAC blurs images of artworks under copyright restriction in ways that hinder transparency to the user. Additionally, because GAC captures some artworks in extremely high-resolution images, there is a potential to print and misuse them for personal or commercial need [113]. This has led to GAC failing to win the trust of some artists to give permission of bringing their art into Google's digital museum, principally because it does not meet owners' requirement R4/owner:know—the ability to know how the digital asset is used—nor custodians' requirements R12/custodian:access and R17/custodian:usage—the abilities to manage access control to the digital assets and metadata, and to monitor and analyse the usage of the assets by third parties.

The flagship project of the European Union, **Europeana**, is another leading service for data management of cultural heritage. Officially released in 2009, Europeana's web portal contains over 58 million digitised cultural heritage records of more than 3,600 institutions across Europe [36]. It offers a single access point to Europe's digital cultural and scientific heritage aggregated from libraries, archives, museums and audio-visual archives. Europeana offers various services for different stakeholders such as citizens, cultural heritage professionals and the creative industry, publishing its aggregated content as Linked Data, supporting requirements 16 and R34/eco:entity. It thoroughly covers end users' requirement R23/enduser:explore, allowing users to explore the digital archive by a plethora of features, including collection topics, type of media, copyright specifications, providing country, language, institution, colour, even image orientation and size [35], and as such allowing them to compare and confront digital material from multiple sources, covering requirement R25/enduser:compare. Furthermore, for citizens, Europeana is especially successful at allowing end users to enhance digital collections, requirement R26/enduser:enhance. For example, in the Europeana Migration campaign [30] over 3000 citizens

contributed to a migration thematic collection by sharing their personal migration stories and accompanying pictures, diaries, videos and letters, and Europeana recently announced the 2020 “Gif It Up” competition for most creative reuse of digitized cultural heritage material [85].

Table 2 provides an overview of the surveyed technologies for data management and the mapping with citizen curation requirements. We observe that none of these technologies meet any owner requirement. Although R14/custodian:monitor and R15/custodian:inappropriate aren’t satisfied by any existing technologies, the majority of the custodian needs could be addressed by data management systems. Moreover, a good coverage of the end user, builder and ecosystem requirements emerges from the analysis.

Table 2. List of technologies for data management and mapping to citizen curation requirements.

	Type	Owner	Custodian	Builder	End user	Eco	Maturity
Omeka S <sup>5</sup>	Sys		6, 7, 8, 9, 10, 11, 12		25	31, 30	Stable
DSpace [100]	Sys		6, 7, 8, 9, 10, 13	18, 22	23, 28	37	Stable
Fedora [83]	Sys		6, 7, 8, 9, 10, 12, 16	18, 22	23, 28	30, 31, 32, 33, 34, 37	Stable
Islandora <sup>6</sup>	Sys		6, 7, 8, 9, 10, 12, 13, 16	18, 22	23, 25, 26, 28	30, 31, 32, 33, 34, 37	Stable
Samvera <sup>7</sup>	Sys		6, 7, 8, 9, 10, 12, 13, 16	18, 22	23, 25, 26, 28	30, 31, 32, 33, 34, 37	Stable
ResearchSpace [77]	Sys		6, 7, 8, 9, 10, 13, 16	18, 22	23, 28	30, 31, 37	Stable
Google A&C [69]	Sys				23, 24, 25, 26, 29		Stable
Europeana [36]	Sys		16		23, 25, 26	34	Stable

#### 4.2 Web technologies for publication and exchange of metadata

We survey Web Technologies and their application for metadata aggregation in cultural heritage, reusing the survey conducted by Freire et al. [38] and complement it with recent advances, namely the Linked.Art project, the W3C ActivityPub, and The Solid Project.

The most established technology for publishing and aggregating digital archives is the **Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH)**<sup>9</sup>. Based on XML technologies and the basics notions of producer and consumer, it’s adoption has been widespread before the advent of the REST approach for developing Web APIs [38]. OAI-PMH implements requirement R8/custodian:publish, R22/builder:standards, R31/eco:consume, R37/eco:extend.

In the context of cultural heritage, a key initiative towards a truly distributed infrastructure is the **International Image Interoperability Framework (IIIF)**<sup>10</sup> focused on supporting the publishing and exchange of high quality images on the Web, supporting authentication, access, and presentation of digital images through dedicated Web APIs. Requirements covered by IIIF are: R8/custodian:publish, R18/builder:interop,

<sup>9</sup><https://www.openarchives.org/pmh/>

<sup>10</sup><https://iiif.io/>

R22/builder:standards, R12/custodian:access, R31/eco:consume, 30, and R34/eco:entity. Also, the standard explicitly supports end-user applications developing browsers and explorators of digital content - requirement R23/enduser:explore.

**Sitemaps** enable webmasters to provide a structured layout of the website content to Web engines, providing a list of URLs associated with an extensible set of metadata in XML. This technologies build on a file specification (Sitemap) and a protocol<sup>11</sup>. Digital libraries on the Web can use Sitemaps to provide a structured catalogue of digital objects within the library [38]. The requirements covered are: R8/custodian:publish, R22/builder:standards, R31/eco:consume, R37/eco:extend.

**ResourceSync** is a standard developed by NISO for synchronizing repositories supporting versioning and notifications of changes [51]. It is based on Sitemaps protocol, extended with Websub for supporting notifications. However, its application in the domain of cultural heritage is rather limited [38]. The citizen curation requirements covered by ResourceSync are: R8/custodian:publish, R22/builder:standards, R31/eco:consume, R37/eco:extend, R32/eco:collect.

**The Open Publication Distributions System (OPDS)** is based originally on the Atom XML syndication format mainly directed to support eBook reading systems, publishers, and providers<sup>12</sup>. The currently developed 2.0 draft adopts JSON-LD and Schema.org and include support for catalogues of images<sup>13</sup>. Overall, the specification supports a subset of the requirements already covered by IIF: R8/custodian:publish R18/builder:interop, R16/custodian:link, R22/builder:standards, R31/eco:consume.

**Webmention** [82] tackles the problem of receiving notifications when other content is published which refer to a document owned by the receiver. When an agent publishes or updates a document, it can send a notification to the target URL mentioned in the document, to notify the target. This is often a blog POST but it can be any type of content. The protocol builds on very basic HTTP techniques, such as the Link header or a HTTP form post. No mechanism is provided to assess trust, therefore, the specification suggests to validate the mention by downloading the source. Notifications do not include data transmission, so the technology needs to be complemented by other means for performing content update. This approach could be used to satisfy requirements: R22/builder:standards, R32/eco:collect.

**WebSub** provides an option for developing connections across distributed systems based on HTTP web hooks [39]. The protocol assumes *hubs* managing the discovery and exchange of content, without requiring any commitment on the format of the document. Content publishers and subscribers interact by the means of *topics*. Agents can discover the hub (or more hubs) managing content of a topic URL, and perform a HTTP POST to subscribe. Publishers of the *topic* notify the hub of changes, which in turns broadcast the message to all subscribers. WebSub covers requirements related to monitoring and tracing content changes, with particular focus on feeds such as Atom or RSS. The specification supports the following requirements: R8/custodian:publish R18/builder:interop R22/builder:standards, R31/eco:consume, and R32/eco:collect.

**Schema.org** [46] is the only technology mentioned by Freire et al. [38] that deals with modelling aspects of metadata aggregation for cultural heritage. In our analysis, its relevance relates to the fact that it represents a set of techniques used for publishing structured metadata. Data is incorporated into HTML pages using various techniques (RDFa, Microformats). In our analysis, we leave out considerations related to knowledge representation but focus on its capacity to develop consensus from the bottom up within communities of practitioners of the Web of data. Cultural heritage institutions can use Schema.org to publish structured metadata on Web pages, which in turns can be scraped by agents to extract structured information. Requirements

<sup>11</sup><https://www.sitemaps.org/protocol.html>

<sup>12</sup>The OpenPub Community. OPDS Catalog 1.1 specification (2011): <https://specs.opds.io/>

<sup>13</sup><https://drafts.opds.io/opds-2.0>



covered are: R8/custodian:publish, R16/custodian:link, R22/builder:standards, R34/eco:entity, and R37/eco:extend.

**The Linked.Art project** follows the steps of IIF with the aim of providing a set of APIs for the exchange of detailed metadata information in cultural heritage, focusing primarily on the art domain<sup>14</sup>. Building on top of Web APIs and the Linked Data format JSON-LD, its vision is to make Linked Data *usable* by relieving the developers from dealing with RDF and heterogeneous vocabularies. Linked.Art can be considered as the development of an ad-hoc API for the exchange of catalogue metadata. Museums can use Linked.Art for publishing content on the Web of Data in a standardized way. Requirements covered are: R8/custodian:publish, R16/custodian:link, R22/builder:standards, and R34/eco:entity.

The **ActivityPub** protocol [114] is a specification for supporting decentralised social networking on the Web, building on top of the ActivityStreams 2.0 data format [101]. ActivityPub provides two APIs, one for client to server communication (for creating, updating and deleting content), and the other for server to server communication, to tackle content distribution. Actors have one inbox and one outbox, and an associated collection of followers, to which servers broadcast new messages. ActivityPub is a promising technology to enable communication of social applications on the Web. Requirements covered are: R18/builder:interop, R37/eco:extend, R16/custodian:link, R22/builder:standards, R24/enduser:share, R34/eco:entity, and R30/eco:link.

**The Solid Project**<sup>15</sup> builds on the RDF technology stack and the REST approach to Web APIs to provide a set of techniques for accessing and manipulating resources as Linked Data. The Solid project aims to develop an ecosystem of specifications and technologies for building decentralised applications, with particular accent on improving privacy and data ownership on the Web. Solid takes the challenge of developing a fully distributed infrastructure on top of the Web of Data. The goals of the project are ambitious and so far, the core specifications cover resource management (Linked Data Platform [1]), content update notification (Linked Data Notifications [17]), identity, and distributed access control (Web Access Control<sup>16</sup>). This technology stack is very promising for citizen curation as it satisfies some key requirements related to copyright, access, distributed identity, and extensibility: R8/custodian:publish, R12/custodian:access, R16/custodian:link, R11/custodian:enable, R6/custodian:manage, R8/custodian:publish, R22/builder:standards, R31/eco:consume, R37/eco:extend, R34/eco:entity, R30/eco:link, R36/eco:manage. However, the Solid Project is under development<sup>17</sup> and, therefore, specifications have different levels of maturity.

Table 3 summarises the result of the survey on the web technologies supporting the citizen curation requirements. Web technologies don't meet any owner requirements and very few end user requirements. These technologies mostly address requirements at ecosystem level. Finally, as for custodian requirements, they are mostly covered by Solid but poorly addressed by other technologies.

### 4.3 Rights data management

A recent survey on open access policy and practice [73] in the GLAM sector shows how more than 70% of the organisations have data that cannot be published openly on the Web. Digital rights management means different things in relation to security enforcement and access control (ACL) (for example, as in [95] and formalisation and reasoning over legal knowledge (terms and conditions, licences) [7, 54]. In this section we focus on technologies for the management of legal knowledge (terms and conditions, licences), considering approaches whose aim is to express and reason upon policies in the meaning of licences, and limiting to the approaches designed to work with the WWW's architecture or principles.

<sup>14</sup>Linked.Art: <https://linked.art/>

<sup>15</sup><https://solidproject.org/>

<sup>16</sup><https://github.com/solid/web-access-control-spec>

<sup>17</sup>The development is open and community-oriented:<https://github.com/solid>.

Table 3. List of Web technologies for metadata exchange and mapping to citizen curation requirements.

	Type	Owner	Custodian	Builder	End user	Eco	Maturity
OAI-PMH <sup>9</sup>	Inf		8	22		31, 37	Stable
IIIF <sup>10</sup>	Inf		8, 12	18, 22	23	31, 34	Stable
Sitemaps <sup>11</sup>	Inf		8,	22	23	31, 37, 32	Stable
ResourceSync [51]	Inf		8	22	23	31, 37, 32	Stable
OPDS <sup>12</sup>	Inf		8, 16	18	23	31, 37, 32	Stable
Webmentation [82]	Inf			22		32	Stable
WebSub [39]	Inf		8	18,22		31, 32	Stable
Schema.org [46]	Inf		8, 16	22		34, 37	Stable
Linked.Art <sup>14</sup>	Inf		8, 16	22		34, 37	Stable
ActivityPub [114]	Inf		16	18, 22	24	30, 34, 37	Stable
Solid <sup>15</sup>	Inf		6, 8, 11, 12, 16	22		31, 30, 34, 37	Under Dev.

Borissova [14], provides a thorough analysis of copyright-related issues raised by cultural heritage digitisation, confirming how the impact of intellectual property is hampering the economic exploitation of cultural heritage, arguing how *"cultural heritage is essential economic resource which uniqueness is national competitive advantage and as such it should be fully industrially utilised but under intellectual property protection."* In particular, cultural heritage is a sector in which intellectual property requires the development of *sui generis* rights, beyond standard definitions [14]. Although initiatives such as rightsstatements.org are influencing organisations to review how rights statements are included in the catalogue metadata, digital archives have limited support for rights data management, often confined to one or two metadata fields (e.g. `dc:terms:rights`) which textual content often includes all sort of information, including non pertinent information such as how the asset was acquired [103]. Instead, rights data management received increasing attention in the Semantic Web research area, in recent years. Therefore, we base our survey on [13] (2007) and [59] (2018), limiting to solutions relevant to rights expression, acquisition, and reasoning, and complement it with works published more recently, which cited the mentioned surveys. Kirrane [59] report on a set of high-level tasks relevant to rights data management, that we report with an alignment with some key requirements for citizen curation (see Table 4).

Our analysis focuses on rights data management, therefore, we do not survey literature about *enforcing* digital rights, excluding techniques to encrypt copyright information in cultural content, for example, DRM approaches, watermarking, and blockchain. In addition, we leave out datasets and domain ontologies, for example, we don't discuss specific work on representing the EU General Data Protection Regulation with RDF/ODRL<sup>18</sup> and only discuss the RDF Licence Database while exploring tools that rely on it. For an overview of the problems of licence compatibility, composition, and propagation, we refer the reader to [29, 42, 111]. Table 5 shows the list of technical solutions surveyed in this section.

**The Open Digital Rights Language (ODRL)** [54] is designed to support a fine-grained expression of rights statements, based on three main components: *permissions*, *prohibitions*, and *duties*. The Semantic Web community developed a number of solutions dealing with policy reasoning on top of the W3C ODRL Ontology [55]. A first order logic semantics for ODRL/XML has been proposed, and used to determine precisely when a permission is implied by a set of ODRL statements, showing that answering such questions is a decidable NP-hard problem [87]. We restrict our analysis to methods based on the Open Digital Rights Language (ODRL) that

<sup>18</sup>For example, within the W3C Data Privacy Vocabularies and Controls Community Group (DPVCG): <https://www.w3.org/community/dpvcg/> (accessed, 9/11/2020).

Table 4. List of tasks initially introduced in the survey [59], aligned with our requirements for citizen curation.

Task	Description	Requirements
Policy selection	Select or compose an appropriate policy for an artifact.	R1/owner:express
Policy communication	Disseminate the policy to (potential) consumers.	R3/owner:negotiate, R19/builder:request, R5/owner:claim, R29/enduser:claim
Monitoring	Monitor the use and distribution of the artifact for policy management.	R4/owner:know, R17/custodian:usage
Policy enforcement	Put mechanisms in place to enforce compliance with the policy.	[Not applicable]
Policy interpretation	Interpret the implications of the policy in their own context.	R21/builder:explain, R20/builder:viewterms
Compatibility testing	Check that the policy is compatible with that of artifacts they are consuming/producing.	R21/builder:explain, R27/enduser:privacy
Usage monitoring	Track usage of the artifact for policy compliance	R17/custodian:usage, R27/enduser:privacy
Validation	Check that usage of the artifact is compliant with the policy	R21/builder:explain, R27/enduser:privacy

Table 5. List of solutions for rights management and mapping to citizen curation requirements.

	Type	Owner	Custodian	Builder	End user	Eco	Maturity
The Open Digital Rights Language (ODRL) [54]	Inf	1, 3, 5	11	19, 20, 22	28, 37		Stable
Licentia[18]	Inf	1		20, 21	28		Stable
Licence Picker [28]	Inf	1		20, 21	28		Stable
CaLi	Inf			20, 21	28		Stable
SPINdle	Inf	3	17				Stable
DALICC	Inf	1, 2, 3, 4		20	27		Stable
PPR	Inf	4	14, 17		28		Stable
PrivOnto	Inf	1		20, 21	28, 27		Stable
SPECIAL	Inf	1, 2, 3, 4			27		Stable
rightsstatements.org	Inf			20, 22			Stable

we consider superseded legacy technologies such as MPEG-21<sup>19</sup>. The RDF Licence Database is an example of a resource built with ODRL to improve communication and explanation of licences by means of a dataset of over 100 licenses [92]. ODRL supports several requirements related to the expression and computation of rights: R1/owner:express, R3/owner:negotiate, R5/owner:claim, R11/custodian:enable, R19/builder:request, R20/builder:viewterms, R22/builder:standards, R28/enduser:copyright, and R37/eco:extend.

<sup>19</sup><https://en.wikipedia.org/wiki/MPEG-21>

The problem of licence identification and selection is an important one. Tools such as TLDRLegal<sup>20</sup>, CC Choose<sup>21</sup>, and ChooseALicense<sup>22</sup> help users to browse licences and select the appropriate one for their resources. However, these do not typically rely on a formal representation of the rights. The RDF Licence Database mentioned before is at the base of the tools **Licentia** [18] and **Licence Picker** [28]. Licentia is a Web system that aims at making it easier for users to select a licence to associate to an asset. Specifically, Licentia aims at supporting producers in understanding license terms, in checking the compatibility of a given licence with the aims of the owner, and to support a graphical visualisation. Licentia can provide support for: R1/owner:express, R20/builder:viewterms, R21/builder:explain, R28/enduser:copyright. **Licence Picker** uses an ontology – LiPio, developed from the RDF Licence Database. LiPio was built applying Formal Concept Analysis (FCA) as a method for clustering licences with respect to their formal specification and develop a workflow based on a set of questions, designed by curating the clusters produced by FCA. The resulting workflow allows to reach a decision by answering 3-5 questions, thus reducing the effort in licence identification<sup>23</sup> [28]. The approach could be applied to cluster custom set of ODRL requirements and provide the basis for an effective integration with user interfaces for the acquisition of rights information. Licence Picker could support: R1/owner:express, R20/builder:viewterms, R21/builder:explain, R28/enduser:copyright. Applying a similar approach, **CaLi** applies Formal Concept Analysis (FCA) for automatically position a given licence over a set of licenses with relation to compatibility and compliance. The system implements the classification technique into a license-based search engine for the Web of Data, supporting requirements R20/builder:viewterms, R21/builder:explain, and R28/enduser:copyright.

Licence compatibility testing is a crucial feature in data management pipelines for citizen curation. **SPINdle** is the reasoner used by Licentia to perform tasks such as licence checking and compatibility [64]. It implements a logic solver able to reason over deontic statements, involving permissions, prohibitions, and duties. This reasoning can support requirements such as negotiation of terms of use: R3/owner:negotiate, R17/custodian:usage. Likewise, the production of derivative data works is a central component of citizen curation. The **Data Licenses Clearance Center (DALICC)** supports legal experts, businesses, and developers in the safe reuse of third party digital assets such as datasets or software [84]. Specifically, DALICC provides support for determining which asset can be shared with whom and under which conditions, thus lowering the burden of rights clearance. The system involves four components: a licence library of ready-made licences, a licence composer that reuses existing licences to create custom ones, a negotiator implementing tasks such as compatibility and conflict resolution, and a licence annotator to support the viewing of human readable rights. This framework supports the requirements: R1/owner:express, R2/owner:rights, R3/owner:negotiate, R4/owner:know, R20/builder:viewterms, and R27/enduser:privacy.

Monitoring usage is important for curators and builders when assisting consumers in understanding the licences and supporting the compliant use of protected resources [59]. The **Policy Propagation Reasoner (PPR)** [29] is based on a methodology for acquiring policy and process knowledge [25] and was evaluated in a complex scenario of smart city data management [24]. The technique involves the application of a dataset of licences (in ODRL) and the acquisition of information about processes reusing protected resources with the Datanode Ontology<sup>24</sup> [23, 27], and supports the analysis on how the output of those processes (derived resources) may be affected but the licences. This approach can contribute on coping with requirements R4/owner:know, R14/custodian:monitor, R17/custodian:usage, and R28/enduser:copyright.

<sup>20</sup><https://tldrlegal.com/>

<sup>21</sup><https://creativecommons.org/choose/>

<sup>22</sup><https://choosealicense.com/>

<sup>23</sup>Licence Picker: <http://data.open.ac.uk/licence-picker>

<sup>24</sup>Datanode: <http://purl.org/datanode/ns/>

The Usable Privacy Policy Project (UPP)<sup>25</sup> developed a corpus of more than 20k privacy policies from Web sites. The corpus is used to develop semi-automatic approaches for analyzing privacy policies including crowdsourcing, natural language processing, and machine learning. Automatic annotating documents about rights can significantly help the fruition of the content by non experts. **PrivOnto** [78] is a semantic framework for the analysis of privacy policies. The system is one result of UPP that focuses on adopting a combination of natural language processing (NLP), privacy preference modeling, crowdsourcing, and UI design to pragmatically support users in making sense of websites' existing terms and conditions with the aim of empowering users to more privacy-aware Web surfing. PrivOnto makes use of Semantic Web technologies to support users, researchers, and regulators in the analysis of privacy policies at scale. This approach is relevant to requirements: R1/owner:express, R20/builder:viewterms, R21/builder:explain, R28/enduser:copyright, and R27/enduser:privacy. The European General Data Protection Regulation (GDPR) calls for technical and organizational measures to support its implementation. The **SPECIAL** H2020 project develops a set of tools for supporting data controllers and processors to automatically check if personal data management and distribution respect the duties set forth in the European regulation. SPECIAL include a policy language to express consent, policies, and regulatory obligations, and reasoning systems for automated compliance checking. Those can be used to demonstrate that data processing setup by controllers or processors does not violate the expressed consent of data subjects, and the related business processes satisfy the requirements of GDPR [12]. SPECIAL tools can be used to support requirement: R1/owner:express, R2/owner:rights, R3/owner:negotiate, R4/owner:know, and R27/enduser:privacy.

**rightsstatements.org** [45] provides twelve standardised rights statements following the Linked Data practice to be used in online cultural heritage, organised in three categories: *In copyright*, *No copyright*, and *Other*. Organisations can use these rights statements to explain users how online cultural heritage works can be reused. However, the purpose of rightsstatements.org it's not to replace licences and terms of use with an elaborate, machine readable representation of rights, but to give a simple and standardised way of explaining the key features of licences. Online digital archives are increasingly adopting rightsstatements.org, although it was observed how the provided options do not fit all the needs of memory institutions [8]. Requirements covered are: R20/builder:viewterms and R22/builder:standards. Crucially, the approach itself is not extendable to other terms of use (R37/eco:extend).

#### 4.4 Crowd-sourcing and cultural heritage

In recent years, there has been an increasing interest in the cultural domain heritage towards crowdsourcing contributions from citizens [33, 52]. Thanks to the availability of participatory digital platforms and the wave of increasing user-generated contents as a result of the Web 2.0 phenomenon, museums have occasionally launched initiatives for crowdsourcing cultural heritage information from users and foster co-curation processes [79]. These include photographic campaigns, e.g. the *Wiki Loves Art* campaign launched in 2008 for adding pictures to Wikipedia public records of artworks and buildings, social tagging campaigns, like the crowdsourced enrichment of the *Steve.Museum's* folksonomy [107], and also campaigns for the co-curation of museum exhibitions, such as the *Click!* exhibition of the museum of Brooklyn [86]. In a few cases, cultural heritage institutions were beneficiary of crowdsourcing events despite not being the promoters, such as the well-known reCAPTCHA security system [112] that allowed to decipher around 440 million words from scanned books.

Over the years special purpose systems and applications have been designed by cultural heritage institutions to systematically collect and preserve user-generated data relevant to their collections. A notable example includes *Accurator* [58], an annotation tool designed by the Rijksmuseum of Amsterdam to get domain experts in several disciplines involved in enriching museum's metadata. However, researchers agree that solutions developed for a single institution turn out being not sustainable or not motivating for users [62]. Therefore new

<sup>25</sup><https://explore.usableprivacy.org/>

general-purpose crowdsourcing infrastructures and systems have been developed in order to support multiple stakeholders' requirements and to allow gathering and sharing distributed contents. To the best of our knowledge, a comprehensive account of experiences in participatory cultural heritage is detailed in [93], and the most updated surveys on systems and infrastructures for crowdsourcing cultural heritage data include [10, 11, 62, 80, 91, 106]. Unfortunately, most of surveyed systems are not maintained anymore (e.g. Storify, SNOOPS, MUSE), therefore we narrow the analysis to systems and infrastructures still available and running (see Table 6). Surveyed technologies span from applications for Smart Cities to tool for collaborative annotation of images, texts, and tagging.

Table 6. List of infrastructures and systems for crowdsourcing and mapping to citizen curation requirements.

	Type	Owner	Custodian	Builder	End user	Eco	Maturity
Culture Gate[62]	Inf		12, 15	18, 22	23, 26, 25, 24		Under dev.
Mirador [108]	Sys		6, 11, 12, 13	18, 22, 19, 20	26, 25	30, 31, 37	Stable
Pybossa [11]	Inf		7, 8, 11, 13, 14		26, 25, 27, 28		Stable
Zooniverse [89, 99]	Sys	2	7, 8, 11, 13, 14		23, 26, 28		Stable

**Culture Gate** Culture gate [62] is an infrastructure for participatory cultural heritage management. The platform allows users to access personalised contents on the basis of their location, and to collect information about artefacts they encountered (relevant to requirements R23/enduser: explore, R26/enduser: enhance). In a second phase cultural heritage professionals check for users' data accuracy (R15/custodian: inappropriate), and information is presents to other groups of users (R25/enduser: compare and R24/enduser: share). A strong emphasis is given to security, forbidding unauthorized access to data for users that do not have the sufficient rights and permissions (R12/custodian: access). It allows third-party software to be integrated and new services to be built on top of it (R22/builder: standards, R18/builder: interop).

**Mirador** Mirador [108] is a system that enables the annotation and comparison (R26/enduser: enhance, R25/enduser: compare) of images served by repositories that support the International Image Interoperability Framework (IIIF) API's (R18/builder: interop, R22/builder: standards, R19/builder: request, R20/builder: viewterms). It supports authentication, and users' annotations are represented and stored according to W3C Web Annotation standards (relevant to custodians' requirements R13/custodian: viewpoints, R6/custodian: manage, R11/custodian: enable, R12/custodian: access). It is extensible, and can be integrated on top of other applications, such as Omeka S and ResearchSpace. More importantly, it allows users to access and annotate at the same time images belonging to several museums, such as the Getty Museum and the British Museum, therefore addressing R30/eco: link, R31/eco: consume, R37/eco: extend.

**Pybossa** Pybossa is a GDPR-compliant infrastructure that allows several types of crowdsourcing tasks (e.g., transcription, tagging) that can be completed by participants as specified in custom application templates, reviewed by experienced participants and/or researchers, and analysed, monitored, and exported (R7/custodian: curate, R8/custodian: publish, R11/custodian: enable, R13/custodian: viewpoints, R14/custodian: monitor). Collected data are released under a CC-BY licence on the website hosting the crowdsourcing initiative (R26/enduser: enhance, R25/enduser: compare, R27/enduser: privacy, R28/enduser: copyright). A notable application of the software is Micropasts [11], a project supported by the British Museum and several museums for the transcription and 3D photo-masking of archeological artefacts.

**CATTI** Another area where crowdsourcing has been used is Handwritten Text Recognition (HTR), that uses the general public to transcribe historical manuscripts. This can be with or without the use of automatic handwritten text recognition and represents a variety of Computer Assisted Transcription of Text Images (CATTI) applications. While many of the CATTI systems were for expert use, some use the Web to make use of crowdsourcing in order

to transcribe the text. Some systems are based on the **Zooniverse** [89, 99] platform and include Shakespeare's World, Ancient Lives, Recital [44], and Transcribe Bentham [19]. Advantages of Zooniverse projects, that we remark are other functionalities as annotating tools, are planned workflows, advance user help systems for the inexperienced and breaking down tasks. Other systems include Tikkoun Soferim [115] and MONK-GIWIS [16]. These systems are of interest for us because they archive cultural heritage artefacts, and make them available for the transcription task and archive the results. In terms of our requirements these are systems that are run by custodians so they meet the following custodian requirements R7/custodian:curate, R8/custodian:publish, R11/custodian:enable, R13/custodian:viewpoints, R14/custodian:monitor, while most systems the custodians are also the owner, so problems of rights are not dealt with, some systems such as Tikkoun Soferim use IIF (mentioned above) to manage image owners requirements R2/owner:rights. For the users they let them explore the documents R23/enduser:explore, and enhance them with the transcription R26/enduser:enhance however most ignore end user's right to the material they create R28/enduser:copyright.

#### 4.5 Social Media and cultural heritage

Most memory institutions tend to have a presence on social media and this tends to be focused on major platforms such as Facebook, Twitter and YouTube. Zafiroopoulos et al. [118], in an analysis of 57 museums, found that 42 had Twitter accounts, 45 had Facebook accounts, and 33 had both. adilla-Meléndez et al. [81], in analysis of the 40 most popular museums online found that their main social media presence was on Facebook (67.5%), Twitter (55%) and YouTube (42.5%). Similarly, an evaluation of social media use by US museums found the most commonly used platforms to be Facebook (94%), Twitter (70%) and YouTube (56%) [37]. Although museums tend to have a presence on these major social media platforms, they are largely used for broadcasting to the public rather than enabling participatory experiences. Schick and Damkjær [94], found that citizen engagement with museums on social media “rarely advances beyond small talk and the content shared lacks any immediately apparent theoretical or cultural importance” (p. 37).

Nevertheless, a recent survey Vassiliadis et al. [109], reviewed 54 papers presenting studies about the use of social media in museums and examined what are the principal axioms museums tackle with social media implementation, and from the survey, emerged four effects: benefits of social media within museums, social media effects on learning process, insights about the use of social media in museums and problems and barriers of social media integration in museums. When considering social media within museums, the authors note that social media may boost dialog, real time communication and engagement, encourage participation learning, encourage cultural consumption beyond the museum boundaries and greatly extend the communication channels the museum may have with visitors. Social media may also be relevant from marketing point of view. When considering the effect on learning, using social media museums can inspire high educational engagement. The insights about the use of social media in museums are not surprising, museums use social media to communicate and engage audience. They use it for awareness, comprehension and engagement of users as efficient social media promotes discussion and boost visitors' loyalty. They note that different social media tools can be used for achieving different goals. With respect to the use of social media, they refer to a study ([37]) that investigated how 315 museums use social media. They observed that museums managers use social media to post reminders (60%), online promotions, announcements (45%), or to expand their brand awareness and reach new visitors (42%) and that only a minority (11%) uses them to create a bidirectional communication. Finally, they discussed the problems and barriers for integrating social media in museums and claim that museums tend to ignore their responsibility to engage audience and ignore the importance of ICT as a tool. They point out that museums lack IT knowledge which may be a reason why they do not exploit its full potential. They also note about ethical practice of social media being a principal issue for museums. From this overview we can conclude that museums do not use social

media at their full potential and that a possible explanation for this might be the lack of personnel's expertise and knowledge about how to use mass social media effectively.

From our survey, the social media that are mostly used by museum professionals are: Facebook, Twitter, Instagram, Pinterest, YouTube, and Flickr. In Table 7, we summarise requirements covered by such platforms. Although we appreciate that there are inherent differences between the services, we can observe how mass social media support similar requirements for citizen curation. They all provide means for end users (including also custodians) to share content, opinions, and viewpoints - which are relevant to requirements R8/custodian:publish, R9/custodian:visibility, R24/enduser:share, R13/custodian:viewpoints, R16/custodian:link, R23/enduser:explore -, to have some form of control on the appropriateness of shared contents (R15/custodian:inappropriate). In addition, they stimulate an economy of systems and tools built on top of ad-hoc Web APIs for interacting with the published content - R18/builder:interop, including permitting third-party applications to record and publish new content (R20/builder:viewterms).

Table 7. Citizen curation requirements satisfied by social media platforms.

	Type	Owner	Custodian	Builder	End user	Eco	Maturity
Facebook	Inf		8, 9, 13, 16, 15, 24	18,20	23		Stable
Twitter	Inf		8, 9, 13, 16, 15, 24	18,20	23		Stable
Instagram	Inf		8, 9, 13, 16, 15, 24	18,20	23		Stable
Pinterest	Inf		8, 9, 13, 16, 15, 24	18,20	23		Stable
Flickr	Inf		8, 9, 13, 16, 15, 24	18,20	23		Stable

#### 4.6 Distributed Online Social Networks

Research on Distributed Online Social Network (DOSN) focuses on connecting persons and their content without a single intermediary organisation that can *observe* content sharing, messaging, and other user activities. Here we focus on analysing how approaches in this area of research may fit some of the requirements of citizen curation, reusing results from a number of surveys in the area: [20] that focuses on approaches based on P2P networks; [104], which analyses aspects of security and privacy; [61] illustrates a number of open challenges for DOSN, for example, in relation to scalability and access control; and [48], which covers in addition issues of data availability and information diffusion. Those surveys should be consulted for a general overview of the field.

There is a variety of paradigms approaching the development of a DOSN, from solutions based on P2P networks [20] including solutions built on Web technologies such as Persona and Diaspora [61]. However, there are several engineering challenges related to content management in DOSN [48]: (1) Data availability/persistence, referring to the fact that nodes of the network may not be always reachable and that the capacity of servers owned by individuals or small organisations may harm content access significantly; (2) Information diffusion, referring to the propagation of content updates in the network; and (3) Privacy, in relation to the problem of providing sophisticated methods for access control management in the distributed system.

In a distributed social network, users do not depend on a central organisation to maintain their data. The lack of a single control point removes several potential privacy breaches. However, without a central authority there is nobody that can guarantee that a trusted source would provide a secured access to protected information and, at the same time, protect users and their content. Examples of platforms for distributed social networks include Mastodon [120], Socialhome<sup>26</sup>, and Hubzilla<sup>27</sup>. The main motivation for adopting DOSN approach for developing

<sup>26</sup>Socialhome: <https://socialhome.network/> (accessed 22/10/2020)

<sup>27</sup>Hubzilla: <https://zotlabs.org/page/hubzilla/hubzilla-project> (accessed 22/10/2019)



citizen curation is the preservation of ownership of the content belonging to museums and cultural heritage providers. In addition, in this way the digital asset stays with its authoritative source, providing a trust-able cultural context, which is typically lost on mass social media.

Table 8 provides a summary of the survey on distributed online social networks in the context of citizen curation. Most of the solutions developed in this area focus on the problem of managing data access from a distribution of nodes, covering some key requirements of custodians, including access control - R12/custodian:access. However, the reference scenario is the one of social media, targeting primarily the exchange of messages between users and the broadcast to "friends" in the network. Crucially, approaches like Diaspora provide a stable framework for developing distributed social media applications allowing to push data to third-party repositories - R36/eco:manage, and relying on a distributed identity management - R33/eco:identity. Mastodon is a distributed platform for microblogging that relies on W3C standards such as ActivityStreams and ActivityPub for interoperability. The project is open source and developers can be exploit it to build distributed social media applications. With the exception of Mastodon, the protocols and data models of DOSN systems are not relying on open standards - R22/builder:standards. All systems seem to be constrained to a set of typical social media data objects and don't support an extensible data model, necessary to cover custodian requirements such as R7/custodian:curate. Unfortunately, this family of tools do not support the delegation of access control, therefore, none of the requirements of the Owner from our workflow model are supported.

Table 8. List of technologies for supporting DOSN and mapping to citizen curation requirements.

	Type	Owner	Custodian	Builder	End user	Eco	Maturity
PeerSon [15]	Inf		8,6,12	18	24	31	POC
LifeSocial.KOM [43]	Inf		8,6,12	18	24	31,37	POC
My3 [76]	Inf		8,6,12	18	23,24	31,33	POC
Gemstone [105]	Inf/Ser				24	31	Stable
SOUP [60]	Inf		8,6,12	18	23,24	31	Stable
Safebook [21]	Inf		8,6,12	18	23,24	31	POC
DiDuSoNet [47]	Inf		12 (limited)	18	23,24	31	POC
Prometheus [63]	Inf		8,6,12	18	23,24	31,33,36	Stable
Vis-a-Vis [97]	Inf/Sys		8,6,12		24	31	POC
Vegas [97]	Inf/Sys		8,6,12 (limited)		24	31	POC
SuperNova [98]	Inf		8,6,12		24	31	POC
Confidant [68]	Inf		8,6,9,12 (limited)		24	31	POC
POSN [34]	Inf				24	31	POC
PrPl [96]	Inf				24	31,33	POC
Diaspora [9]	Inf/Ser		8,6,9,12		24	31,33,36	Stable
Mastodon <sup>28</sup>	Inf/Sys		8,6,9, 16,12		24,22	31,33,36	Stable

## 5 DISCUSSION, CHALLENGES, AND OPPORTUNITIES

In order to answer the question **What are the requirements of citizen curation?** we derived 37 requirements from the scenario and the model outlined in Section 3, targeting needs of data owners (R1-R5), custodians (R6-R17), builders (R18-R22), end users (R23-R29), and non-functional requirements (R30-R37). To answer the question

**What are the technologies that contribute to an infrastructure for citizen curation?** in Section 4 we surveyed 55 systems, infrastructures, and services from six research areas (see Table 1).

In this section we explore the solution space to answer the following questions: *To what extent existing systems inter-operate? Are there any missing critical components in the state of the art? What types of connections need to be established? What approach should we take in order to fill the gaps?* We analyse the data collected in the survey<sup>29</sup> and we summarise the results of the survey in three figures.

In Figure 3, we show what technologies cover requirements relevant to a certain role. Technologies are visually grouped by research area (see bottom lines), so as to highlight which roles are mostly represented in the area. For instance, Dalicc and SPECIAL cover the higher number of requirements relevant to owners. Almost all technologies falling under the research area "Rights Management" address owners' requirements. Compared to other technologies, owners are mostly (and only) represented in this research area.

In Figure 4, we show which research areas cover specific requirements. Values are normalised in percentage, so as to characterise the interconnection between areas and their contribution to the solution. For instance, the second column represents the requirement R2/owner:rights, which is addressed by, respectively, two technologies falling under Rights Management area and one technology falling under Crowdsourcing area. That is, technologies falling under Rights Management area are more likely (70%) to fulfil the requirement rather than technologies falling under Crowdsourcing area (30%).

Lastly, in Table 9 we show the average distribution (again, in percentage) of requirements (grouped by role) covered by technologies (grouped by research area), which tells us, on average, to what extent technologies falling under a certain research area contribute to cover the requirements of a user role. For instance, in the first row we observe how owners' requirements are met by areas. Most requirements are fulfilled by technologies falling under "Rights Management", which in average cover 30% of requirements for that role.

All together, the three figures allow us to characterise research areas through the lenses of technologies (Fig. 3), requirements (Fig. 4), and roles (Table 9).

First, we explore our findings in relation to the ability of the research areas to support requirements of the various roles. According to our initial assumption (Table 1), we expect owners' requirements (R1-R5) to be satisfied by technologies developed in the context of research on Rights management. This is confirmed by Figure 4 and Figure 3, particularly by resources like ODRL, Dalicc, and Special. A minor contribution is given by crowdsourcing tools. However, by looking at Table 9, we understand that technologies developed in Rights management cover (on average) 32% of owner's requirements. Likewise, custodians' requirements (R6-R17) are expected to be tackled by Data management, Web, and Rights management technologies. In Figure 4 we can observe that technologies for Data Management contribute to almost all requirements (with an average of 42.9%, see Table 9), except Google Arts and Culture and Europeana that are mostly devoted to end users' needs. Instead, Web technologies are far less present, and Rights Management tackles only a few requirements. Surprisingly, Social Media and Crowdsourcing contribute on average (respectively) to satisfy 50%, 33.3% of custodians' requirements. Builders' requirements (R18-R22) are expected to be addressed by Data management, Web technologies, and Rights management. While this seems to be true for Rights Management (that contribute to four out of five requirements), Web technologies and Data Management contribute to satisfy only two out of five requirements (see Figure 4). Noticeable is instead the contribution of technologies for Crowdsourcing and Social Media. Nonetheless, the average contribution of each research area to tackle builders' requirements is similar, as shown in Table 9, with an exceptional higher contribution from Social Media. End users' requirements (R23-R29) are expected to be the focus of Crowdsourcing, Rights management, and Data Management. While the assumption is confirmed, again we find a higher contribution from Social Media, which in average contribute

<sup>29</sup>Data is available online for reproducibility purposes. Daga, Enrico et al. (2021): Survey on technologies for citizen curation. figshare. Dataset. <https://doi.org/10.6084/m9.figshare.14578887.v1>

to 42.9% to satisfy users' requirements (Table 9). Finally, non-functional requirements (R30-R37) are mainly tackled by Web technologies and tools for Data management (as expected). Among the others, we notice that technologies for Distributed Social Network offer interesting contribution to the field, by tackling in average 20.3% of ecosystem requirements.

To answer the question **To what extent existing systems support cooperation?** we evaluate how much the technologies surveyed cover requirements belonging to multiple roles. In order to do that, we produce a correlation matrix from the point of view of the roles, in Table 11. The data shows a strong correlation between systems supporting custodians and end-users, demonstrating how crowd-sourcing initiatives and curatorial activities are tightly connected in the state of the art. In addition, from the correlation between *custodian* and *ecosystem* we can see that technologies supporting non-functional requirements are somehow supporting also the needs to memory institutions, and vice-versa. However, the analysis makes emerge also weaknesses of the current state of affairs. The needs of builders and end-users seem to be somehow orthogonal, which is not entirely surprising considering that the builder essentially act as a mediator between users and custodians. The problem is that builders seem to receive limited support in general, confirming the idea that the cultural industry does not have a structured, resilient approach to building infrastructure for citizen curation. Builders do not have means for discovering protected content, requesting access and negotiating terms of use. Instead, experiences are occasional and typically result in isolated, stand-alone applications. This conclusion relates to the analysis of data management systems (Section 4.1), although technologies and standards for supporting expression and negotiation of rights, and distribution of identity and control exist in areas such as Web technologies (Section 4.2) or distributed online social networks (Section 4.6) - Solid being representative of the first area, and Mastodon of the second. Solutions for rights expression, negotiation, and management exist (Section 4.3) but they are not used by the sector. The *rightsstatements.org* initiative is too limited in scope to be considered relevant for the issues raised by citizen curation. Crowd-sourcing initiatives (Section 4.4) are confined to end-to-end, silos applications and do not cover issues such as control of rights and ownership, and cross-platform identity management. Social media (Section 4.5) have limited support for cooperation especially for the reason that content access and transfer is problematic (for example, Twitter terms of use don't allow to copy the content to third-party databases). Owners' requirements are covered by technologies that are very different from the ones that typically are used for supporting curators and end-users. Specifically, we can conclude that copyright and licence data management are poorly supported by systems targeted to custodians, builders, and end-users.

To answer the question **Are there any missing critical components in the state of the art?** we focus on the requirements that have received minimal support by the technologies included in the survey, and observe how they related to the various research areas. In particular, we compare our expectations on research areas contributing to requirements (Table 1, and Figure 4). We already observed how owners' requirements are poorly supported – for example, artists or users of social media. In particular, owners do not have agency in the current state of affairs, and cannot directly grant or revoke rights – R2/owner:rights, don't have means to know how a digital asset is used – R4/owner:know, and data management systems don't offer support for claiming ownership of the content – R5/owner:claim. These issues are similar to requirements R19/builder:request, R20/builder:viewterms, R21/builder:explain, and R29/enduser:claim, which are not supported by current solutions in the data management area. Although knowledge representation and reasoning technologies to support these tasks exist (from the Rights data management research area), there is still work to be done in understanding how they can be applied to cultural heritage. Another neglected area of requirements relates to the monitoring of content in digital archives of custodians. Although we did not explore how content monitoring is performed on mass social media – for which we refer the reader to [4], it is evident how cultural heritage archives do not have means for dealing with the heterogeneity of issues related to user generated content – R14/custodian:monitor, R15/custodian:inappropriate. However, Web technologies already offer basic components to allow external agents to react to the production and update of content in third-party repositories (for example, ActivityPub).




Fig. 3. Coverage of requirements grouped by roles in surveyed technologies.

Table 9. Distribution of requirements grouped by role (in percentage) over technologies grouped by research area

	Data Mngmt	Web tech.	Rights Mngmt	Crowd-S	DOSN	Social M.
owner	0	0	32	5	0	0
custodian	42.9	13.6	3.3	33.3	21.4	50
builder	28.6	27.3	28	30	10	40
enduser	38.8	2.6	12.9	42.9	18.75	42.9
ecosystem	33.9	30.7	2.5	12.5	20.3	0

Intelligent agents could monitor user generated content and raise alerts when potential issues are identified, for example, with relation to privacy or copyright violations – R17/custodian: usage, by combining technologies for rights data management. These shortcomings are also reflected in relation to ecosystem requirements. Although R30/eco:link and R34/eco:entity are naturally supported by any system based on Linked Data principles, issues such as distribution of identity have limited support – R33/eco:identity. In particular, data management systems may have the capacity of delegating the authentication to an external system (for example, with OAuth 2.0 [50]) but does not prevent the disclosure of the identity of the end-user (the mail address). However, citizen curation expects end-users to interact with systems provided by third-parties, which act as mediators between the user and the organisation (custodian). These mediators will need to handle user identity and deal with commitments related to privacy law (GDPR), avoiding to propagate personal identifiable information to remote parties. Although the identity may be unknown to the organisation collecting and curating the content, users may still want to be able to have control over the final content. The problem cannot be addressed until requirements R33/eco:identity and R35/eco:agency will get attention by the research community on data infrastructures for cultural heritage. The Solid Project and research in DOSN may be able to offer a solution.

To answer the question **What types of connections need to be established?** we looked at the correlation between research areas (Table 10) and at the correlation between roles (Table 11). It is evident how data management technologies should evolve to incorporate solutions for rights data management. Interestingly, technologies for Distributed Online Social Network have a high correlation with data management solutions, meaning that they can solve similar problems. Those approaches may provide the basis for as truly distributed




Fig. 4. Coverage of requirements in surveyed research areas (in percentage).

content management infrastructure, possibly based on Linked Data principles. Rights management are still isolated from the other research areas relevant to citizen curation. As discussed in the previous sections, both data and identities need to be available on the Web but access must be controlled by content owners or delegated to a trusted agent (custodian). Identities should be linked to digital assets through metadata annotations, and these metadata need also to be accessed with appropriate restrictions. These prerequisites make the context fundamentally different from the one of (Linked) Open Data, where personal identifiable information can be removed through anonymisation techniques and one can expect to expose the same content to any user without distinctions. However, it is an open question how an infrastructure for supporting Linked Non-Open Data would look like.

A specific issue for the creation of citizen curation platforms from existing systems and frameworks concerns the lack of dedicated workflows for the type of user-generated content produced by citizen curation processes. Thanks to their architectural modularity and the extendibility of their representation schemas, most systems, especially in the data management area, have in principle the capability to represent user-generated content, but they don't acknowledge this type of content as part of their workflows. In the current situation, any attempt to use these systems to include user generated content would fall short of creating the appropriate paths for handling them according to the requirements. Current workflows, in fact, mostly rely on an unidirectional path from ingestion to fruition, with user responses not being reintroduced in the system as first class citizens. The flexibility of access tools provided by most systems, which rely on effective indexing modules, is generally intended for the end user, and it is not available to create curation paths depending on content types and features (in other terms, they are not ready to enable the *scripting* of activities for generating and managing user-generated contents). Also, extending the current solutions to create citizen curation systems may not be feasible for small organisations, that cannot afford a similar effort (in terms of know-how, costs, staff, infrastructures) needed to make significant development work on top of their current solutions. We can conclude from the analysis that the role of mediators will be crucial in supporting the heterogeneity and diversity of organisations involved in citizen curation projects, providing specific services (e.g. licence clearance or monitoring inappropriate content) and flexible, cost-effective platforms to design and curate interactions, processes, and data. However, such mediators will need an underlying infrastructure able to deal with distribution of content, identity, and control, following the path traced by the Solid Project and the DOSN experiences.

Our discussion so far focused on the limitations of current technologies with relation to the challenges raised by citizen curation, while we would like to conclude this survey by tackling the question **What approach**

Table 10. Correlation matrix of research areas

-	Data Management	Web tech.	Rights Management	Crowd-S	DOSN	Social M.
Data Management	1					
Web technologies	0.39	1				
Rights Management	-0.8	-0.33	1			
Crowdsourcing	0.76	$-7.8 \cdot 10^{-2}$	-0.3	1		
DOSN	0.95	0.38	-0.34	0.58	1	
Social Media	0.67	$-6.4 \cdot 10^{-2}$	-0.19	0.93	0.45	1

Table 11. Correlation matrix of roles

-	owner	custodian	builder	enduser	ecosystem
owner	1				
custodian	-0.66	1			
builder	0.28	0.13	1		
enduser	-0.24	0.85	0.25	1	
ecosystem	-0.83	0.55	$-4.8 \cdot 10^{-2}$	$3.2 \cdot 10^{-2}$	1

**should we take in order to fill the gaps?** Some of those challenges above described were inherently part of Semantic Web research in the last twenty years. Cultural heritage is a field where semantic technologies have already been introduced and employed in the past two decades. Semantic web technologies, in particular, could support many of the requirements reviewed in this article. For what concerns the semantic data management, however, it emerges that little advancements have been made from the agenda settled in 2012 [110]. Specifically, still insufficient efforts have been made for what concerns the availability of integrated knowledge systems, the validation of the extended data models provided by each cultural institution, the capability of handling uncertain reasoning, the multilingualism of the exposed cultural repositories etc. In general, a crucial element to improve would concern the integration of the existing technologies with the standard workflow operation of Cultural Experts and Information Professionals (IPs) in Libraries, Archives and Museums (LAMs) [74]. Finally, and consistently with the requirements provided by the citizens curation ecosystem, the crucial issues concerning the ownership, permission of use, trust, and copyrights issues have only been recently sketched. The above mentioned Solid Project represents the front-runner initiative in the field to answer to such requirements. However, the project is very ambitious and the technologies involved have different levels of maturity. Specifically, the community is still far from tackling important issues such as the one of information discovery and diffusion - issues which a technological infrastructure aiming at supporting citizen curation should necessarily be able to cope with. However, we have seen how research on Distributed Online Social Network provided solution to securely broadcast protected content. The Mastodon project is based on Web technologies (ActivityPub) and at the same time offers a social media platform based on content distribution. However, the type of content and interactions supported seems limited. Overall, there is definitely an opportunity in attempting to combine Linked Data and the Solid principles with the approaches studied in the area of DOSN to try to fill these gaps. Finally, blockchain technologies are receiving increasing attention from the art market community for its ability to certify transactions and therefore support the tracing of authenticity of artworks in auctions [116]. For an analysis of blockchain technologies and their relevance for digital archives, we refer the reader to [32]. Blockchain could support the task of tracing digital assets use - R14/custodian:monitor and certifying negotiated terms of use R4/owner:know. However, the analysis of the impact of citizen curation on technologies for the enforcement

of access control and digital rights, including methods to validate processes and policies, is an important future work.

## 6 CONCLUSIONS

In this article we studied the novel paradigm of citizen curation from the point of view of an ecosystem of technologies for data management. We identified four roles in citizen curation, the owner, the custodian, the builder, and the end-user, and contributed an analysis of their requirements. We have shown how existing research can contribute significantly in facing the challenges raised by citizen curation. However, there are still some important gaps that need to be filled. These are related to the limited support for distribution, primarily in relation to ownership, identities, workflows, and monitoring. The Solid Project and DOSN research may indicate a path towards possible solutions.

How do we curate and share our cultural heritage? Raymond [88] and Roued-Cunliffe and Copeland [93], makes us reflect on two possible *models*: a cathedral and a bazaar. The cathedral model is a carefully constructed edifice, well defined, with clear boundaries and with a unified design. Instead, the bazaar model is open, it is an ongoing structure with a wide variety of approaches. The existing platforms seem oriented to building cathedrals, even when they aim at opening the content to others (CultureGate, ResearchSpace). Citizen curation, instead, requires a *bazaar* of solutions to work together in an open-ended ecosystem. The SPICE project aims at building an infrastructure to allow cultural heritage assets as well as citizens' opinions, responses, and memories to be shared within safe channels. This type of infrastructure can provide support for developing applications that exploit digital content among different organisations, connecting users' contributions from a multiplicity of engagement approaches and supporting services that span different platforms while preserving the privacy, ownership, and fair use of the resources involved. This "participatory" approach fits in to the reappraisal of who is the *expert* in cultural heritage [93], with the museum curator no longer necessarily being the sole expert, thus requiring methods and technologies that are accessible to many different types of people with various skill levels (which can be bridged) and differing goals and motives. It is an open question whether the cultural heritage sector and industry will be able to take the challenge.

## ACKNOWLEDGMENTS

The research leading to this publication has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreements "SPICE - Social Cohesion, Participation, and Inclusion through Cultural Engagement" (Grant Agreement N. 870811), <https://spice-h2020.eu>, and "Polifonia: a digital harmoniser of musical cultural heritage" (Grant Agreement N. 101004746), <https://polifonia-project.eu>. The publication reflects the author's views. The Research Executive Agency (REA) is not liable for any use that may be made of the information contained therein.

## REFERENCES

- [1] John Arwe, Steve Speicher, and Ashok Malhotra. 2015. *Linked Data Platform 1.0*. W3C Recommendation. W3C. <https://www.w3.org/TR/2015/REC-ldp-20150226/>.
- [2] Stavroula Bampatzia, Angeliki Antoniou, George Lepouras, Costas Vasilakis, and Manolis Wallace. 2016. Using social media to stimulate history reflection in cultural heritage. In *2016 11th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP)*. IEEE, IEEE, Piraeus Conference Centre, 12-14 Katouni Str. 546 26 Thessaloniki, 89–92.
- [3] Jo Bates, Yu-Wei Lin, and Paula Goodale. 2016. Data journeys: Capturing the socio-material constitution of data objects and flows. *Big Data & Society* 3, 2 (2016), 2053951716654502.
- [4] Bogdan Batrinca and Philip C Treleaven. 2015. Social media analytics: a survey of techniques, tools and platforms. *Ai & Society* 30, 1 (2015), 89–116.
- [5] Peter Bazalgette. 2017. *The empathy instinct: How to create a more civil society*. Hachette UK, UK.

- [6] Agiatis Benardou, Erik Champion, Costis Dallas, and Lorna Hughes. 2017. *Cultural heritage infrastructures in digital humanities*. Routledge, Abingdon, Oxfordshire, UK. <https://doi.org/10.4324/9781315575278>
- [7] V Richard Benjamins, Pompeu Casanovas, Joost Breuker, and Aldo Gangemi. 2005. Law and the semantic web, an introduction. In *Law and the Semantic Web*. Springer, 1–17.
- [8] Sara R Benson and Hannah Stitzlein. 2020. Copyright and Digital Collections: A Data-Driven Roadmap for Rights Statement Success. *College & Research Libraries* 81, 5 (2020), 753.
- [9] Ames Bielenberg, Lara Helm, Anthony Gentilucci, Dan Stefanescu, and Honggang Zhang. 2012. The growth of diaspora—a decentralized online social network in the wild. In *2012 proceedings IEEE INFOCOM workshops*. IEEE, IEEE, 13–18.
- [10] Tobias Blanke, Conny Kristel, and Laurent Romary. 2015. Crowds for clouds: recent trends in humanities research infrastructures. *arXiv preprint arXiv:1601.00533* (2015).
- [11] Chiara Bonacchi, Andrew Bevan, Adi Keinan-Schoonbaert, Daniel Pett, and Jennifer Wexler. 2019. Participation in heritage crowdsourcing. *Museum Management and Curatorship* 34, 2 (2019), 166–182.
- [12] Piero A Bonatti, Sabrina Kirrane, Iliana M Petrova, and Luigi Sauro. 2020. Machine Understandable Policies and GDPR Compliance Checking. *arXiv preprint arXiv:2001.08930* (2020).
- [13] Piero A. Bonatti and Daniel Olmedilla. 2007. Rule-based Policy Representation and Reasoning for the Semantic Web. In *Proceedings of the Third International Summer School Conference on Reasoning Web (Dresden, Germany) (RW'07)*. Springer-Verlag, Berlin, Heidelberg, 240–268. <http://dl.acm.org/citation.cfm?id=2391482.2391488>
- [14] Vladia Borissova. 2018. Cultural heritage digitization and related intellectual property issues. *Journal of Cultural Heritage* 34 (2018), 145–150.
- [15] Sonja Buchegger, Doris Schiöberg, Le-Hung Vu, and Anwitaman Datta. 2009. PeerSoN: P2P social networking: early experiences and insights. In *Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*. ACM, Nuremberg, 46–52.
- [16] Anna Caceres, Andreas Weber, and Lambert Schomaker. 2020. MONK in Practice: Indexing Heterogeneous Handwritten Collections. In *7th Digital Humanities Benelux 2020*, Vol. 7. DH Benelux, Leiden, Netherlands, –.
- [17] Sarven Capadislis and Amy Guy. 2017. *Linked Data Notifications*. W3C Recommendation. W3C. <https://www.w3.org/TR/2017/REC-ldn-20170502/>.
- [18] Cristian Cardellino, Serena Villata, Fabien Gandon, Guido Governatori, Brian Lam, and Antonino Rotolo. 2014. Licentia: a Tool for Supporting Users in Data Licensing on the Web of Data. In *Proceedings of the ISWC 2014 Posters & Demonstrations Track, a track within the 13th International Semantic Web Conference (ISWC 2014)*, Matthew Horridge, Marco Rospocher, and Jacco van Ossenbruggen (Eds.). Riva del Garda, Italy.
- [19] Tim Causer and Melissa Terras. 2014. Crowdsourcing Bentham: beyond the traditional boundaries of academic history. *International Journal of Humanities and Arts Computing* 8, 1 (2014), 46–64.
- [20] Shihabur Rahman Chowdhury, Arup Raton Roy, Maheen Shaikh, and Khuzaima Daudjee. 2015. A taxonomy of decentralized online social networks. *Peer-to-Peer Networking and Applications* 8, 3 (2015), 367–383.
- [21] Leucio Antonio Cutillo, Refik Molva, and Thorsten Strufe. 2009. Safebook: A privacy-preserving online social network leveraging on real-life trust. *IEEE Communications Magazine* 47, 12 (2009), 94–101.
- [22] Paulo Pinheiro da Silva, Silva Deborah, Deborah L McGuinness, and Rob Mccool. 2003. Knowledge provenance infrastructure. In *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering*. IEEE.
- [23] Enrico Daga. 2018. The Datanode Ontology v0.5. <https://doi.org/10.5281/zenodo.1293998> <http://purl.org/datanode/0.5/ns/>.
- [24] Enrico Daga, Mathieu d’Aquin, Alessandro Adamou, and Enrico Motta. 2016. Addressing exploitability of smart city data. In *2016 IEEE International Smart Cities Conference (ISC2)*. IEEE, Trento, Italy, 1–6.
- [25] Enrico Daga, Mathieu d’Aquin, Aldo Gangemi, and Enrico Motta. 2015. Propagation of policies in rich data flows. In *Proceedings of the 8th International Conference on Knowledge Capture*. 1–8.
- [26] Enrico Daga, Mathieu d’Aquin, and Enrico Motta. 2017. Propagating data policies: a user study. In *Proceedings of the Knowledge Capture Conference*. 1–8.
- [27] Enrico Daga, Mathieu d’Aquin, Aldo Gangemi, and Enrico Motta. 2014. Describing semantic web applications through relations between data nodes. *Technical Report kmi-14-05, Knowledge Media Institute, The Open University, Walton Hall, Milton Keynes* (2014).
- [28] Enrico Daga, Mathieu d’Aquin, Enrico Motta, and Aldo Gangemi. 2015. A Bottom-Up Approach for Licences Classification and Selection. In *The Semantic Web: ESWC 2015 Satellite Events*. Springer, Springer, Portorož, Slovenia. <https://link.springer.com/book/10.1007%2F978-3-319-25639-9>
- [29] Enrico Daga, Aldo Gangemi, and Enrico Motta. 2018. Reasoning with data flows and policy propagation rules. *Semantic Web* 9, 2 (2018), 163–183.
- [30] Beth Daley. 2019. Curated by Europeana - and the people of Europe. <https://pro.europeana.eu/post/curated-by-europeana-and-the-people-of-europe>
- [31] Conseil de l’Europe. 2006. *Council of Europe Framework Convention on the Value of Cultural Heritage for Society: Faro, 27. X. 2005*. Éd. du Conseil de l’Europe, Faro.



- [32] Angelica Lo Duca, Clara Bacciu, and Andrea Marchetti. 2020. The Use of Blockchain for Digital Archives: a comparison between Ethereum and Hyperledger (AIUCD 2019). *Umanistica Digitale* 4, 8 (2020).
- [33] Stuart Dunn and Mark Hedges. 2014. How the crowd can surprise us: Humanities crowdsourcing and the creation of knowledge. *Crowdsourcing our cultural heritage* (2014), 231–246.
- [34] Esra Erdin, Eric Klukovich, Gurhan Gunduz, and Mehmet Hadi Gunes. 2015. Posn: A personal online social network. In *IFIP International Information Security and Privacy Conference*. Springer, 51–66.
- [35] Europeana. [n.d.]. Search. <https://www.europeana.eu/en/search>
- [36] Europeana. 2020. The Europeana platform: Europe’s digital cultural collection for responsible, accessible, sustainable and innovative tourism. <https://ec.europa.eu/digital-single-market/en/europeana-platform-europes-digital-cultural-collection-responsible-accessible-sustainable-and>
- [37] Adrienne Fletcher and Moon J. Lee. 2012. Current social media uses and evaluations in American museums. *Museum Management and Curatorship* 27, 5 (2012), 505–521. <https://doi.org/10.1080/09647775.2012.738136> arXiv:<https://doi.org/10.1080/09647775.2012.738136>
- [38] Nuno Freire, Glen Robson, Antoine Isaac, Hugo Manguinhas, and John B Howard. 2017. Web Technologies: A Survey of Their Applicability to Metadata Aggregation in Cultural Heritage.. In *ELPUB*. 235–244.
- [39] Julien Genestoux and Aaron Parecki. 2018. *WebSub*. W3C Recommendation. W3C. <https://www.w3.org/TR/2018/REC-websub-20180123/>.
- [40] Elisa Giaccardi. 2012. *Heritage and social media: Understanding heritage in a participatory culture*. Routledge, Abingdon, Oxfordshire, UK.
- [41] Tarleton Gillespie. 2010. The politics of ‘platforms’. *New media & society* 12, 3 (2010), 347–364.
- [42] Guido Governatori, Antonino Rotolo, Serena Villata, and Fabien Gandon. 2013. One License to Compose Them All. In *The Semantic Web–ISWC 2013*. Springer, Sydney, Australia, 151–166.
- [43] Kalman Graffi, Christian Gross, Dominik Stingl, Daniel Hartung, Aleksandra Kovacevic, and Ralf Steinmetz. 2011. LifeSocial. KOM: A secure and P2P-based solution for online social networks. In *2011 IEEE Consumer Communications and Networking Conference (CCNC)*. IEEE, IEEE, 554–558.
- [44] Adeline Granet, Benjamin Hervy, Geoffrey Roman-Jimenez, Marouane Hachicha, Emmanuel Morin, Harold Mouchère, Solen Quiniou, Guillaume Raschia, Françoise Rubellin, and Christian Viard-Gaudin. 2018. Crowdsourcing-based Annotation of the Accounting Registers of the Italian Comedy. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA), Miyazaki, Japan, 446,450. <https://www.aclweb.org/anthology/L18-1069>
- [45] International Rights Statements Working Group et al. 2015. Rightsstatements.org white paper: Recommendations for standardized international rights statements. Retrieved from [rightsstatements.org/files/160208recommendations\\_for\\_standardized\\_international\\_rights\\_statements\\_v11](https://rightsstatements.org/files/160208recommendations_for_standardized_international_rights_statements_v11) (2015).
- [46] Ramanathan V Guha, Dan Brickley, and Steve Macbeth. 2016. Schema.org: evolution of structured data on the web. *Commun. ACM* 59, 2 (2016), 44–51.
- [47] Barbara Guidi, Tobias Amft, Andrea De Salve, Kalman Graffi, and Laura Ricci. 2016. DiDuSoNet: A P2P architecture for distributed Dunbar-based social networks. *Peer-to-Peer Networking and Applications* 9, 6 (2016), 1177–1194.
- [48] Barbara Guidi, Marco Conti, Andrea Passarella, and Laura Ricci. 2018. Managing social contents in decentralized online social networks: a survey. *Online Social Networks and Media* 7 (2018), 12–29.
- [49] Juliet L Hardesty. 2014. Exhibiting library collections online: Omeka in context. *New Library World* (2014).
- [50] Dick Hardt et al. 2012. *The OAuth 2.0 authorization framework*. Technical Report. RFC 6749, October.
- [51] Bernhard Haslhofer, Simeon Warner, Carl Lagoze, Martin Klein, Robert Sanderson, Michael L Nelson, and Herbert Van de Sompel. 2013. ResourceSync: leveraging sitemaps for resource synchronization. In *Proceedings of the 22nd International Conference on World Wide Web*. 11–14.
- [52] Mark Hedges and Stuart Dunn. 2017. *Academic crowdsourcing in the humanities: Crowds, communities and Co-production*. Chandos Publishing.
- [53] Amanda Hill, Mark Kretzschmar, David Morton, and Sara Raffel. 2018. Chapter Six “Eenie Meenie Miney Mose”: Using Experimental Citizen Curating To Engage Visitors With Racial Ephemera. *Florida Studies Review* (2018), 62.
- [54] Renato Iannella. 2001. Open digital rights management. In *World Wide Web Consortium (W3C) DRM Workshop*. W3C, Sophia-Antipolis, France.
- [55] Renato Iannella, Michael Steidl, Stuart Myles, and Victor Rodriguez-Doncel. 2017. *ODRL Version 2.2 Ontology*. Technical Report. W3C. <https://www.w3.org/ns/odrl/2/>
- [56] ICOM. 2020. *Museum Definition*. <https://icom.museum/en/resources/standards-guidelines/museum-definition>
- [57] Bente Jensen. 2013. Instagram as cultural heritage: User participation, historical documentation, and curating in Museums and archives through social media. In *2013 Digital Heritage International Congress (DigitalHeritage)*, Vol. 2. IEEE, 311–314.
- [58] Lizzy Jongma and Chris Dijkshoorn. 2016. Accurator: Enriching Collections with Expert Knowledge from the Crowd. In *Proceedings of the museums and the web conference*.

- [59] Sabrina Kirrane, Serena Villata, and Mathieu d'Aquin. 2018. Privacy, security and policies: A review of problems and solutions with semantic web technologies. *Semantic Web* 9, 2 (2018), 153–161.
- [60] David Koll, Jun Li, and Xiaoming Fu. 2014. Soup: an online social network by the people, for the people. In *Proceedings of the 15th International Middleware Conference*. 193–204.
- [61] David Koll, Jun Li, and Xiaoming Fu. 2017. The good left undone: Advances and challenges in decentralizing online social networks. *Computer Communications* 108 (2017), 36–51.
- [62] Zois Koukopoulos, Dimitrios Koukopoulos, and Jason J Jung. 2017. A trustworthy multimedia participatory platform for cultural heritage management in smart city environments. *Multimedia Tools and Applications* 76, 24 (2017), 25943–25981.
- [63] Nicolas Kourtellis, Joshua Finnis, Paul Anderson, Jeremy Blackburn, Cristian Borcea, and Adriana Iamnitchi. 2010. Prometheus: User-controlled p2p social data management for socially-aware applications. In *ACM/IFIP/USENIX International Conference on Distributed Systems Platforms and Open Distributed Processing*. Springer, Springer, 212–231.
- [64] Ho-Pun Lam and Guido Governatori. 2009. The making of SPINdle. In *International Workshop on Rules and Rule Markup Languages for the Semantic Web*. Springer, Springer, Las Vegas, Nevada (USA), 315–322.
- [65] Luciana Lazerretti, Andrea Sartori, and Niccolò Innocenti. 2015. Museums and social media: the case of the Museum of Natural History of Florence. *International Review on Public and Nonprofit Marketing* 12, 3 (2015), 267–283.
- [66] Jin Woo Lee, Yikyung Kim, and Soo Hee Lee. 2019. Digital Museum and User Experience: The Case of Google Art & Culture, In Proceedings of the 25th International Symposium on Electronic Art. *International Symposium on Electronic Art*. <https://kar.kent.ac.uk/75700/>
- [67] Jasmine Li. 2020. Omeka Classic vs. Omeka. net. *Emerging Library & Information Perspectives* 3, 1 (2020), 232–236.
- [68] Dongtao Liu, Amre Shakimov, Ramón Cáceres, Alexander Varshavsky, and Landon P Cox. 2011. Confidant: protecting OSN data without locking it up. In *ACM/IFIP/USENIX International Conference on Distributed Systems Platforms and Open Distributed Processing*. Springer, Springer, 61–80.
- [69] Google LLC. Accessed October 25, 2020. *About Google Arts and Culture*. <https://about.artsandculture.google.com/experience/>
- [70] Peter Lyman and Brewster Kahle. 1998. Archiving digital cultural artifacts. *D-lib Magazine* 4, 7 (1998).
- [71] Deborah Maron and Melanie Feinberg. 2018. What does it mean to adopt a metadata standard? A case study of Omeka and the Dublin Core. *Journal of Documentation* (2018).
- [72] Barry J Mauer. 2017. The Citizen Curating Project Confronts the Pulse Nightclub Shooting. *The St. John's University Humanities Review* 14 (2017). Issue 1. <http://stjenglish.com/wp-content/uploads/St.-JohnUniversity-Humanities-Review-14.2Spring-2017-1.pdf>
- [73] D McCarthy and A Wallace. 2018. Survey of GLAM open access policy and practice. *Copyright Cortex*. 18 (2018), 2020. <http://bit.ly/OpenGLAMsurvey>
- [74] Lucy McKenna, Christophe Debruyne, and Declan O'Sullivan. 2018. Understanding the position of information professionals with regards to linked data: a survey of libraries, archives and museums. In *Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries*. ACM/IEEE, 7–16.
- [75] Heather Moqtaderi. 2019. Citizen Curators: Crowdsourcing to Bridge the Academic/Public Divide. *University Museums and Collections Journal* 11, 2 (2019), 204–210.
- [76] Rammohan Narendula, Thanasis G Papaioannou, and Karl Aberer. 2012. A decentralized online social network with efficient user-driven replication. In *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*. IEEE, IEEE, 166–175.
- [77] Dominic Oldman and Diana Tanase. 2018. Reshaping the Knowledge Graph by Connecting Researchers, Data and Practices in ResearchSpace. In *The Semantic Web - ISWC 2018 - 17th International Semantic Web Conference, Monterey, CA, USA, October 8-12, 2018, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 11137)*, Denny Vrandečić, Kalina Bontcheva, Mari Carmen Suárez-Figueroa, Valentina Presutti, Irene Celino, Marta Sabou, Lucie-Aimée Kaffee, and Elena Simperl (Eds.). Springer, 325–340. [https://doi.org/10.1007/978-3-030-00668-6\\_20](https://doi.org/10.1007/978-3-030-00668-6_20)
- [78] Alessandro Oltramari, Dhivya Piraviperumal, Florian Schaub, Shomir Wilson, Sushain Cherivirala, Thomas B Norton, N Cameron Russell, Peter Story, Joel Reidenberg, and Norman Sadeh. 2018. PrivOnto: A semantic framework for the analysis of privacy policies. *Semantic Web* 9, 2 (2018), 185–203.
- [79] Johan Oomen and Lora Aroyo. 2011. Crowdsourcing in the cultural heritage domain: opportunities and challenges. In *Proceedings of the 5th International Conference on Communities and Technologies*. 138–149.
- [80] Trevor Owens. 2013. Digital cultural heritage and the crowd. *Curator: The Museum Journal* 56, 1 (2013), 121–130.
- [81] Antonio Padilla-Meléndez and Ana Rosa del Águila-Obra. 2013. Web and social media usage by museums: Online value creation. *International journal of information management* 33, 5 (2013), 892–898.
- [82] Aaron Parecki. 2017. *Webmention*. W3C Recommendation. W3C. <https://www.w3.org/TR/2017/REC-webmention-20170112/>.
- [83] Sandra Payette and Carl Lagoze. 1998. Flexible and Extensible Digital Object and Repository Architecture (FEDORA). In *Proceedings of the European Conference on Research and Advanced Technology for Digital Libraries, Second European Conference (ECDL 1998)*. Springer, 41–59. [https://doi.org/10.1007/3-540-49653-X\\_4](https://doi.org/10.1007/3-540-49653-X_4)

- [84] Tassilo Pellegrini, Victor Mireles, Simon Steyskal, Oleksandra Panasiuk, Anna Fensel, and Sabrina Kirrane. 2018. Automated rights clearance using semantic web technologies: The DALICC framework. In *Semantic Applications*. Springer, 203–218.
- [85] Europeana Pro. 2020. GIF IT UP. <https://pro.europeana.eu/page/gif-it-up>
- [86] Laurel Ptak. 2009. CLICK!: A CROWD-CURATED EXHIBITION. *Aperture* 194 (2009), 12–12. <http://www.jstor.org/stable/24473378>
- [87] Riccardo Pucella and Vicky Weissman. 2006. A formal foundation for ODRL. *arXiv preprint cs/0601085* (2006).
- [88] Eric Raymond. 1999. The cathedral and the bazaar. *Knowledge, Technology & Policy* 12, 3 (1999), 23–49.
- [89] Jason Reed, M Jordan Raddick, Andrea Lardner, and Karen Carney. 2013. An exploratory factor analysis of motivations for participating in Zooniverse, a collection of virtual citizen science projects. In *2013 46th Hawaii International Conference on System Sciences*. IEEE, Honolulu, Hawaii, 610–619.
- [90] Peter Ride. 2013. Creating#citizencurators: putting Twitter into museum showcases.. In *Proceedings of the 19th International Symposium of Electronic Art, ISEA2013*. ISEA International, Sydney, Australia.
- [91] Mia Ridge. 2013. From tagging to theorizing: deepening engagement with cultural heritage through crowdsourcing. *Curator: The Museum Journal* 56, 4 (2013), 435–450.
- [92] Victor Rodriguez-Doncel, Serena Villata, and Asunción Gómez-Pérez. 2014. A dataset of RDF licenses.. In *JURIX*. 187–188.
- [93] Henriette Roued-Cunliffe and Andrea Copeland. 2017. *Participatory heritage*. Facet Publishing.
- [94] Lea Schick and Katrine Damkjær. 2010. Can you be friends with an art museum? Rethinking the art museum through Facebook. In *The Digital Turn*. Peter Lang, Tartu, 36. <https://www.peterlang.com/view/9783653023251/9783653023251.00010.xml>
- [95] David Scott and Richard Sharp. 2002. Abstracting application-level web security. In *Proceedings of the 11th international conference on World Wide Web*. ACM, ACM, 396–407.
- [96] Seok-Won Seong, Jiwon Seo, Matthew Nasielski, Debangsu Sengupta, Sudheendra Hangal, Seng Keat Teh, Ruven Chu, Ben Dodson, and Monica S Lam. 2010. PrPI: a decentralized social networking infrastructure. In *Proceedings of the 1st ACM Workshop on Mobile Cloud Computing & Services: Social Networks and Beyond*. ACM, San Francisco, California, USA., 1–8. <https://dl.acm.org/doi/proceedings/10.1145/1810931>
- [97] Amre Shakimov, Harold Lim, Ramón Cáceres, Landon P Cox, Kevin Li, Dongtao Liu, and Alexander Varshavsky. 2011. Vis-a-vis: Privacy-preserving online social networking via virtual individual servers. In *2011 Third International Conference on Communication Systems and Networks (COMSNETS 2011)*. IEEE, IEEE, Bangalore, India., 1–10.
- [98] Rajesh Sharma and Anwitaman Datta. 2012. Supernova: Super-peers based architecture for decentralized online social networks. In *2012 Fourth International Conference on Communication Systems and Networks (COMSNETS 2012)*. IEEE, IEEE, Bangalore, India, 1–10.
- [99] Robert Simpson, Kevin R Page, and David De Roure. 2014. Zooniverse: observing the world’s largest citizen science platform. In *Proceedings of the 23rd international conference on world wide web* (Seoul, Korea). ACM, New York, NY, USA, 1049–1054.
- [100] MacKenzie Smith, Mary R. Barton, Margret Branschofsky, Greg McClellan, Julie Harford Walker, Michael J. Bass, David Stuve, and Robert Tansley. 2003. DSpace: An Open Source Dynamic Digital Repository. *D Lib Mag.* 9, 1 (2003). <https://doi.org/10.1045/january2003-smith>
- [101] James Snell and Evan Prodmou. 2017. *Activity Streams 2.0*. W3C Recommendation. W3C. <https://www.w3.org/TR/2017/REC-activitystreams-core-20170523/>.
- [102] Jocelyn Spence, Benjamin Bedwell, Michelle Coleman, Steve Benford, Boriana N Koleva, Matt Adams, Ju Row Farr, Nick Tandavanitj, and Anders Sundnes Løvlie. 2019. Seeing with new Eyes: Designing for in-the-wild museum gifting. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow, Scotland, UK, 1–13. <https://dl.acm.org/doi/proceedings/10.1145/3290605>
- [103] Hannah Stitzlein, Myung-Ja K Han, and Sara R Benson. 2018. Unraveling challenges: Rights statements in digital cultural heritage collections. *Journal of Library Metadata* 18, 3-4 (2018), 135–150.
- [104] Sanaz Taheri-Boshrooyeh, Alptekin Küpçü, and Öznur Özkasap. 2015. Security and privacy of distributed online social networks. In *2015 IEEE 35th International Conference on Distributed Computing Systems Workshops*. IEEE, IEEE, Columbus, OH, USA, 112–119.
- [105] Florian Tegeler, David Koll, and Xiaoming Fu. 2011. Gemstone: empowering decentralized social networking with high data availability. In *2011 IEEE Global Telecommunications Conference-GLOBECOM 2011*. IEEE, IEEE, 1–6.
- [106] Melissa Terras. 2015. *Crowdsourcing in the Digital Humanities*. John Wiley & Sons, Ltd, Chapter 29, 420–438. <https://doi.org/10.1002/9781118680605.ch29> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118680605.ch29>
- [107] Jennifer Trant, Bruce Wyman, et al. 2006. Investigating social tagging and folksonomy in art museums with steve.museum. In *Collaborative Web Tagging Workshop at WWW2006*. Edinburgh, Scotland.
- [108] Joris Van Zundert. 2018. On not writing a review about Mirador: Mirador, IIF, and the epistemological gains of distributed digital scholarly resources. *Digital Medievalist* 11, 1 (2018).
- [109] Chris A Vassiliadis and Zoe-Charis Belenioti. 2017. Museums & cultural heritage via social media: an integrated literature review. *Tourismos* 12, 3 (2017).
- [110] Konstantinos N Vavliakis, Georgios Th Karagiannis, and Pericles A Mitkas. 2012. Semantic Web in cultural heritage after 2020. In *Proceedings of the 11th International Semantic Web Conference (ISWC)* (Boston, MA, USA). 11–15.
- [111] Serena Villata and Fabien Gandon. 2012. Licenses compatibility and composition in the Web of data. In *Proceedings of the Third International Conference on Consuming Linked Data-Volume 905*. CEUR-WS.org, CEUR-WS, 124–135.

- [112] Luis Von Ahn, Benjamin Maurer, Colin McMillen, David Abraham, and Manuel Blum. 2008. recaptcha: Human-based character recognition via web security measures. *Science* 321, 5895 (2008), 1465–1468.
- [113] Ratri Wahyuningtyas. 2017. Eliminating boundaries in learning culture through technology: a review of Google Arts and Culture. In *The 10th International Conference of the Faculty of Language and Arts, Universitas Kristen Satya Wacana*. 179.
- [114] Christopher Webber and Jessica Tallon. 2018. *ActivityPub*. W3C Recommendation. W3C. <https://www.w3.org/TR/2018/REC-activitypub-20180123/>.
- [115] Alan J Wecker, Uri Schor, Dror Elovits, Daniel Stoekl Ben Ezra, Tsvi Kuflik, Moshe Lavee, Vered Raziel-Kretzmer, Avigail Ohali, and Lily Signoret. 2019. Tikkoun sofrim: A webapp for personalization and adaptation of crowdsourcing transcriptions. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*. 109–110.
- [116] Amy Whitaker. 2019. Art and blockchain: A primer, history, and taxonomy of blockchain use cases in the arts. *Artivate* 8, 2 (2019), 21–46.
- [117] Annie Wu, Santi Thompson, Rachel Vacek, Sean Watkins, and Andrew Weidner. 2016. Hitting the road towards a greater digital destination: Evaluating and testing DAMS at the University of Houston Libraries. (2016).
- [118] Kostas Zafropoulos, Vasiliki Vrana, and Konstantinos Antoniadis. 2015. Use of twitter and Facebook by top European museums. *Journal of Tourism, Heritage & Services Marketing* (2015).
- [119] Jamil Zaki. 2019. *The war for kindness: Building empathy in a fractured world*. Broadway Books.
- [120] Matteo Zignani, Sabrina Gaito, and Gian Paolo Rossi. 2018. Follow the “mastodon”: Structure and evolution of a decentralized online social network. In *ICWSM*. AAAI Press, 541–550.