



The case for the Humanities Citation Index (HuCI): a citation index by the humanities, for the humanities

Giovanni Colavizza¹ · Silvio Peroni^{2,3} · Matteo Romanello⁴

Received: 1 October 2021 / Revised: 14 May 2022 / Accepted: 15 May 2022
© The Author(s) 2022

Abstract

Citation indexes are by now part of the research infrastructure in use by most scientists: a necessary tool in order to cope with the increasing amounts of scientific literature being published. Commercial citation indexes are designed for the sciences and have uneven coverage and unsatisfactory characteristics for humanities scholars, while no comprehensive citation index is published by a public organisation. We argue that an open citation index for the humanities is desirable, for four reasons: it would greatly improve and accelerate the retrieval of sources, it would offer a way to interlink collections across repositories (such as archives and libraries), it would foster the adoption of metadata standards and best practices by all stakeholders (including publishers) and it would contribute research data to fields such as bibliometrics and science studies. We also suggest that the citation index should be informed by a set of requirements relevant to the humanities. We discuss four such requirements: source coverage must be comprehensive, including books and citations to primary sources; there needs to be chronological depth, as scholarship in the humanities remains relevant over time; the index should be collection driven, leveraging the accumulated thematic collections of specialised research libraries; and it should be rich in context in order to allow for the qualification of each citation, for example, by providing citation excerpts. We detail the fit-for-purpose research infrastructure which can make the Humanities Citation Index a reality. Ultimately, we argue that a citation index for the humanities can be created by humanists, via a collaborative, distributed and open effort.

Keywords Citation indexing · Open citations · Bibliographic information retrieval · Open research infrastructure

Giovanni Colavizza, Silvio Peroni and Matteo Romanello have contributed equally to this work.

✉ Giovanni Colavizza
g.colavizza@uva.nl
Silvio Peroni
silvio.peroni@unibo.it
Matteo Romanello
matteo.romanello@unil.ch

- ¹ Institute for Logic, Language and Computation (ILLC), University of Amsterdam, Science Park 107, 1098 XG Amsterdam, The Netherlands
- ² Research Centre for Open Scholarly Metadata, Department of Classical Philology and Italian Studies, University of Bologna, via Zamboni 32, 40126 Bologna, Italy
- ³ Digital Humanities Advanced Research Centre (/DH.arc), Department of Classical Philology and Italian Studies, University of Bologna, via Zamboni 32, 40126 Bologna, Italy
- ⁴ Institute of Classical Studies and Archaeology, University of Lausanne, 1015 Lausanne, Switzerland

1 Introduction

Citation indexes are by now part of the research infrastructure in use by most scientists: a necessary tool in order to cope with the increasing amounts of scientific literature being published. However, existing commercial citation indexes are designed for the sciences and have uneven coverage and unsatisfactory characteristics for humanities¹ scholars. This situation has both discouraged the usage of citation indexes and hindered bibliometric studies of humanities disciplines.

The creation of a citation index for the humanities may well appear as a daunting task due to several characteristics of this field, such as its fragmentation into several sub-disciplines, the common practice of publishing research in languages other than English, as well as the amount of scholarship from past centuries that is still waiting to be digitised.

¹ Throughout this paper, we use the term *humanities* as a shorthand for Arts & Humanities (A&H). To a degree, the Social Sciences are also concerned.

Notwithstanding these challenges, we argue that the creation of such an index can be highly beneficial to humanities scholars for, at least, the following reasons. Firstly, humanities scholars have long been relying on information seeking behaviours that leverage citations and reference lists for the discovery of relevant publications—a strategy that citation indexes are designed to support and facilitate. Secondly, a comprehensive citation index for the humanities will be a valuable source of data for researchers willing to conduct bibliometric studies of the humanities. Lastly, capturing the wealth of references to primary and secondary sources contained in humanities literature will allow to create links between archives, galleries, libraries and museums where digitised copies of these sources can increasingly be found.

Before continuing with this paper, we introduce key terminology related to citation indexing that will be used throughout this paper, adopting the definitions from [1]. These are: bibliographic entity, bibliographic resource and bibliographic citation. A **bibliographic entity** is any entity which can be part of the bibliographic metadata of a bibliographic artifact: it can be a person, an article, an identifier for a particular entity (e.g. a DOI), a particular role held by a person (e.g. being an author) in the context of defining another entity (e.g. a journal article) and so forth. A **bibliographic resource** is a kind of bibliographic entity that can cite or be cited by other bibliographic resources (e.g. a journal article), or that contains other resources (e.g. a journal). A **bibliographic citation** is another kind of bibliographic entity: a conceptual directional link from a citing bibliographic resource to a cited bibliographic resource. The citation data defining a particular citation must include the representation of the conceptual directional link of the citation and the basic metadata of the involved bibliographic resources, that is to say, sufficient information to create or retrieve textual bibliographic references for each of the bibliographic resources. Following [2], we say that a bibliographic citation is an open citation when the citation data needed to define it are compliant with the following principles: structured, separate, open, identifiable, available.

The remaining of this paper is organised as follows. In Sect. 2, we discuss previous work on analysing the behaviour of humanities scholars in relation to information retrieval. We also present the main limitations of existing citation indexes, seen from the perspective of the humanities, and outline the main obstacle that citation indexing has faced in this area. In Sect. 3, we argue for the need of a Humanities Citation Index (HuCI from now onwards) and in Sect. 4 we present what we believe are the essential characteristics that such an index should have. We then propose a possible implementation of HuCI, based on a federated and distributed research infrastructure (Sect. 5). We conclude with some considerations on how HuCI relates to recent efforts to create open infrastructures for research.

2 Related work

2.1 On scholarly information retrieval in the A&H

The needs and behaviours of humanities scholars in terms of information seeking have been an active area of study especially in the field of Library and Information Science (LIS), where research on this topic started in the 1980s and early 1990s [3–5]. For a thorough review of the early literature on this topic, see [6][p. 2198] and [7][pp. 19–21]. Determining the information needs and behaviours of humanities scholars was essential for librarians in order to support scholars in their research by devising new library systems or by improving the guidelines for abstracting publications to cater for the specific needs of humanities scholars [8]. What emerges from this literature are also the key strategies for finding bibliographic information that characterise humanities scholarship. Firstly, scholars use proper names extensively when searching as compared with scholars in other disciplines [9–11]. Secondly, a prominent behaviour among humanities scholars is to search for bibliographic information by browsing [4, 10, 12]. A typical example is browsing books in the stacks or shelves of a library. What characterises browsing as opposed to a targeted search is that it favours the serendipitous discovery of relevant information: the physical proximity of books on library shelves, which is related to their subject classification, may in some cases transcend the boundaries of subjects. Finally, a third prominent search strategy is the already mentioned citation chaining with its two variants of backward and forward chaining [4, 13]. The former consists of starting from one publication—the seed document—and then following up the references it contains in order to expand the initial search and to discover other related publications. The latter consists of starting from a seed document and then finding which other publications cite it. Moreover, an empirical study of the information seeking strategies of humanities scholars reports that searching and browsing proved to be rather ineffective strategies for locating information and that citation chaining was the most common behavioural pattern [13][pp. 227–228].

2.2 Citation indexing and the humanities

Citation indexing is commonplace for Science, Technology, Engineering and Mathematics (STEM) literature. Mainstream indexes such as Google Scholar, the Web of Science, Scopus, Dimensions or Semantic Scholar are largely capable of indexing most citations accurately. To be sure, their coverage is still uneven and far from uniform [14, 15]. One of the critical problems which are left open is the uneven coverage of different disciplines, with humanities disciplines usually faring worse than most [16]. Several reasons for this state of affairs have been individuated, which can be grouped into two

categories: intrinsic factors, which depend on the literature itself, and extrinsic factors, which depend on the information environment where citation mining is performed [17].

Intrinsic factors which act as obstacles to citation indexing in the humanities include the more limited availability of born digital or digitised publications, a higher variety of languages and publication venues in use, the practice to publish monographs, complex referencing practices and motivations which limit their automatic processing. These topics have been amply discussed in the literature [18–22]. *Extrinsic factors* have been less the focus of previous work and include, instead, the variety and fragmentation of catalogs, information systems and other sources of unique identifiers and authoritative metadata. These issues are well known more generally in the Galleries, Libraries, Archives and Museums (GLAM) sector. A recent study on metadata aggregation highlights several characteristics of this landscape, among which these fall within what we here refer as extrinsic factors [23]:

- Each GLAM sub-domain (libraries, archives and museums) applies its specific resource description practices and data models.
- All sub-domains embrace the adoption and definition of standards-based solutions addressing description of resources, but to different extents.
- Interoperability of systems and data is scarce across sub-domains, but it is somewhat more common within each sub-domain, at the national and the international levels.

As a consequence of the limitations enacted by both intrinsic and extrinsic factors, it is more difficult to comprehensively index the humanities via citations, a condition that limited the use of quantitative bibliometric methods in this area [24], despite clear progress over recent time [25,26]. The lack of a comprehensive and reliable citation index remains a known and open problem in the humanities [27–29]. Our contribution proposes a way forward which mainly addresses the obstacles posed by extrinsic factors, and is true to the way the humanities communicate research and retrieve scholarly information.

3 The need for a citation index for the humanities

Scholarship in the humanities rests on solid traditions, most crucially developed in the archives, libraries and information studies communities. It is thus worth asking the question: **why do we need a citation index for the humanities?** We advance four motivations: to dramatically improve current scholars' information retrieval capabilities; to interlink presently siloed GLAM information systems; to foster best

practices in terms of referencing and metadata; to provide research data for bibliometrics and science studies.

3.1 Improve scholarly information retrieval

From an information retrieval point of view, citation indexes seem to be the natural evolution of disciplinary and thematic bibliographies (e.g. the *Annual Bibliography of English Language and Literature*² or *L'Année Philologique*³), which are widely used by scholars across the humanities to conduct literature search. A citation index, in fact, can be seen a bibliography whose entries are linked with one another depending on the citations that are found in the full-text of the catalogued publications. Moreover, thematic bibliographies such as the *World Shakespeare bibliography*⁴ or the *International Dante Bibliography*⁵ often provide users with the ability to search for publications related to specific literary works—a functionality that could also be provided by a citation index which captures references to primary sources.

Despite the existence of bibliographies and bibliographic databases, humanities scholars cannot yet fully rely on citation indexes when searching for secondary literature, nor to keep up to date with recent developments (e.g. via citation alerts). As we highlighted above, it is the limited coverage of existing citation indexes more than any intrinsic limitation that has been the decisive factor in discouraging their more systematic adoption in retrieval practices. This need not be a sealed fate. Assuming sufficient coverage, in both quality and quantity, a citation index for the humanities can first and foremost serve the same information retrieval needs these tools provide for in the sciences since decades. It is likely that a non-negligible fraction of humanities scholars already uses services such as Google Scholar and Google Books [30], even in the absence of comprehensive evidence on their coverage and reliability.

Furthermore, a variable yet non-negligible amount of references in the humanities is given to primary sources, such as archival documents or literary works [31]. There has never been a way to count and retrieve all references to a given primary source without painstaking manual work. Knowledge about primary sources, in terms of their existence, location and means of access, takes up a substantial amount of time and training in the humanities, sometimes becoming all too treasured. In principle, both primary and secondary sources should be indexed in the humanities citation index. This will allow anyone to immediately gauge which sources have been used together, where and by whom. In practice, several open

² <https://www.mla.org/Publications/MLA-International-Bibliography>.

³ <https://about.brepolis.net/lannee-philologique-aph>.

⁴ <https://www.worldshakesbib.org>.

⁵ <https://bibliografia.dantesca.it>.

challenges will need to be overcome first, including programmatic access to uniform GLAM metadata.

3.2 Interlink GLAM collections via citations

GLAM information ecosystems often exist in isolated silos: metadata and data are largely made accessible by the specific institution that creates and curates them. Notable exceptions exist, for example national library catalogs and projects such as Europeana. Nevertheless, to the best of our knowledge, no encompassing information retrieval infrastructure exists spanning across GLAM institutional categories, for example, interlinking libraries (L) with archives (A). Citation links extracted from scholarly literature can do just that.

The literature in the humanities in fact contains a wealth of references to primary sources, accumulated over centuries of scholarly work. Within the scope of one project alone, some of the authors were able to extract nearly 700,000 references to primary sources from approximately 1900 books and 5500 journal articles (Venice Scholar Index⁶ [32]). Citation links connect secondary literature, hence library catalogs, with information systems of archives, galleries and museums. They also connect archives, galleries and museums directly by virtue of co-citation relationships (i.e. two resources are connected if they are cited together by a third one). These links effectively constitute a dormant virtual information system which awaits to be digitally materialised. By so doing, a significant acceleration and democratisation to the access of primary sources can be realised, contributing to a broader scholarly and public engagement with these collections as it is currently the case.

Digitally materialising citation links would create incentives to make GLAM information retrieval and research infrastructure increasingly more interoperable and interdependent, to the great benefit of the research community. Citation indexing requires publication data and metadata, which must be made available by publishers and GLAM institutions. We argue that once the benefits of citation indexing will have been made tangible to a sufficient degree, this will create a positive feedback loop for all stakeholders to gradually improve on their practices in order to make citation indexing increasingly easier and to a large extent automatic.

3.3 Improve current practices

The automatic extraction and indexing of structured information, such as bibliographic citations, typically require a high degree of openness and standardisation in the ecosystem it happens in. Citation indexing requires open, standardised and programmatically accessible metadata about primary and secondary sources alike, as well as access to the full text of

scholarly publications. It also benefits from a high degree of uniformity in the referencing practices of authors, which makes reference parsing all the more feasible. Yet, all this is costly, hard and time-consuming. For all stakeholders to strive to higher openness, standardisation and accessibility, we require a positive incentive. We argue that citation indexing, once it reaches a certain threshold, actually provides for one: if a community starts using citation indexes for information retrieval, being indexed increasingly becomes a necessity; hence, related investments will be made.

Citation indexing starts with authors. Referencing practices, sometimes less than uniform and coherent, pose a significant challenge to the automatic extraction of citations (e.g. [20]). Yet, once references become data, and their value as links is immediately made tangible via citation indexing, authors might have more incentives to make their referencing practices syntactically and stylistically more uniform in view of improving their harvesting and correct indexing.

A similar point in case can be made for publishers. On the one hand, proof-checking work can make sure to provide for uniform references with sufficient information for their indexing, similarly to what is provided by several scientific publishers. On the other hand, and more importantly, publishers could sign up (and effectively contribute) to the Crossref and OpenCitations initiatives, making their metadata and citation data available. The existence of a citation index for the humanities should foster participation in such initiatives. Failing that, or considering the backlog of already published publications (especially if printed), GLAM institutions themselves can take a leading role, as we discuss below.

The positive incentive to expose open, standardised and programmatically accessible metadata provided by the citation index will also apply to GLAM institutions, once the benefits of interlinked collections and increased searchability will become apparent. A crucial challenge for us will be to reach a critical mass of citation data to provide for an indispensable service to a sizable share of the research community and, at the same time, initiate the positive incentive for all stakeholders.

3.4 Research data for bibliometrics and science studies

It is well known that the humanities are significantly understudied by the bibliometrics and quantitative science studies community, largely because of the lack of citation data [33]. This has several consequences, among which the separation of qualitative studies on the humanities from analyses grounded in (bigger) data [34]. Furthermore, it also causes a widespread science-as-the-norm/humanities-as-an-exception mindset in bibliometrics and research evaluation as a whole, as if it were the case that citations cannot be used

⁶ <https://venicescholar.dhlab.epfl.ch>.

to study the humanities. To be sure, indexing citations in the humanities is challenging, yet it would allow the bibliometrics and quantitative science studies communities to finally approach the humanities on equal ground with respect to the sciences. The proposed citation index for the humanities can radically alter this state of affairs. First of all, a bibliometrics for the humanities grounded in data as well as theory could finally be developed, in full recognition of the specificities of the humanities [26]. Secondly, citation data in the humanities are very rich, if we consider the varied publication typologies, languages, primary and secondary sources that come into play. As a consequence, citation data from the humanities will require novel methods and approaches that might not only provide insights into these data, but as well inform further developments when applied to citation data from the sciences. The HuCI can essentially put an end to the age of the so-called “non-bibliometric” humanities.

Citation data, once available, have been used for research evaluation. Indicators such as citation counts or the H-index are widespread and have been amply discussed by the bibliometrics community [35]. Recently, public efforts have been made to call for a redress and improvement in the use of citation-based indicators [36].⁷ It will be likely unavoidable to face similar discussions if and when the HuCI materialises. We believe these worries should not prevent it from happening, for the very reasons we just detailed. Furthermore, HuCI could provide for an opportunity to rethink the way we use citation-based indicators in research evaluation. The humanities have a long-lasting tradition of peer review assessment which, when mixed with situated and contextualised metrics (which in turn need not be just citation-based), has the potential to inform research evaluation in the sciences too.

4 The characteristics of a citation index for the A&H

Having clarified why we believe a citation index for the humanities is motivated, we detail here four requirements we propose it should have. These are: comprehensive source coverage and chronological depth, rich information provided to contextualise citations, and a growth strategy driven by institutional collections. We note that we intend these requirements as something to aspire to: they represent end goals more than necessary conditions to begin with.

4.1 Source coverage

Scholars in the humanities use a complementary variety of publication typologies, such as monographs, journal arti-

cles and contributions in edited volumes. Journal articles, the main focus of existing commercial citation indexes, in general account for a small fraction of the output in all the humanities [31]. We thus argue that the first requirement of a citation index for the humanities is complete coverage in terms of publication typologies. A related requirement, or pain point, is multilingualism. Scholarly literature in national languages abounds in most of the humanities, yet this variety is not often captured by digital resources. A case in point is the situation of Classics: 75% of Classics publications contained in JSTOR are written in English, while the language of publications reviewed in *L'Année Philologique* (APh, the most important bibliography in this field) is much more evenly distributed between English, German, Italian and French. In fact, Scheidel [37] reports that, of the publications reviewed by APh in 1992, 30% were written in English, roughly 25% in Italian, 20% in French and 20% in German. Ideally, language should not be a source of bias in the citation index.⁸

As we anticipated above, the second requirement we put forth is the full indexing of citations to primary sources. Interestingly, this requirement compels a discussion of citation granularity: what is the object of a reference which should be considered in a citation index? Typically, for secondary literature we use the level of the work in FRBR terms [39]. Hence, citations are accumulated for, say, a journal article aggregating over all its expressions (e.g. in pre-print and printed versions) or for a book over all its editions, excluding those with major revisions that justify calling it a new work. For primary sources, we typically consider unique items (the lowest FRBR level), for example archival documents or unique artworks. We consider instead works when, say, dealing with critical editions of a classic author, where the editing activity is considered scholarly and the source is printed into editions. All this to say that the choice of the citation aggregation object is far from straightforward for primary sources, and a good rule of thumb is that further aggregation is always possible, while disaggregation can be more difficult to undo. Hence, we recommend lower FRBR citation aggregation levels when in doubt.

4.2 Chronological depth

The humanities are known to publish at a relatively slower pace than other sciences and to keep citing older relevant literature (e.g. [20,40]). This has two consequences for the citation index: first, and foremost, it is crucial to index older literature as well, spanning back ideally to when systematic scholarly referencing became commonplace [41]. Secondly, and this is not a requirement but an opportunity

⁷ Also see the San Francisco Declaration on Research Assessment (DORA): <https://sfдора.org/read>.

⁸ Promoting measures against language bias in the context of research assessment is one of the three key recommendations made by the Helsinki Initiative [38].

we highlight, digitising and making openly available old and out-of-copyright literature, in conjunction with its indexing via citations, would constitute a great service to scholars. It would not only improve the use of such literature, but open up opportunities to study the history of scholarship in the humanities at unprecedented scale and comprehensiveness.

4.3 Rich in context

Previous work has elucidated how the citation semantics in the humanities tend to be rich and varied [40]. This is of crucial importance when using citations for information retrieval: is a citation supportive or dismissing? Is it contextual, perfunctory or does it substantially underpin an argument? The citation index we propose will need to make every effort possible to offer its users all the means necessary to appreciate and understand every citation link. This is mainly done by providing relevant context, within the bounds of existing copyrights.

Citation contexts are the excerpts of text preceding and following a citation. The most common context, in this sense, is the sentence where a citation is made. Nevertheless, a context can cover any relevant span, e.g. a few sentences or a whole paragraph. Another source of contextual information is given by proximal co-citations: which other sources are cited with the one under consideration, within the same citing publication? Lastly, providing the exact details of the citation, such as the page number it refers to, also helps to specify its scope. It is possible to see citations and their contexts in aggregate, from the point of view of either the citing or cited sources. This is the case when we consider, for example, all the other sources a given source is co-cited with. It is also possible to consider every citation as situated in a quite specific location of a publication. For example, by considering co-cited sources within the same paragraph of a well-defined citing publication. Both views, the aggregate and the detail, provide for relevant contextual information for a scholar to interpret citation links, and to use them for information retrieval.

4.4 Collection driven

We conclude this section not by discussing a requirement, but by suggesting a growth strategy for HuCI. Mainstream citation indexes convey the impression, and sometimes the illusion, of comprehensive coverage. Only when we are able to trust a citation index in this sense, we, as scholars, can rely on it for our work. If a citation index is manifestly incomplete, and especially if what is missing is unknown or hard to qualify, it will be difficult for it to succeed. Given the daunting task we have set ourselves to with the humanities citation index and the stated requirements, we also need a reasonable growth strategy. Our proposal is to be collection driven. That is to say, we recommend to index topically coherent

batches of scholarly literature, by leveraging the specialised collections of research libraries.

In our previous work on the historiography of Venice, we faced the task of defining the limits of what pertains to this topic and what can be left outside. By relying on a set of finding aids—library catalogs, bibliographies, shelving strategies and specialised collections—we were able to create a coherent citation corpus [42]. We suggest here that this approach can make HuCI scale, one themed collection at the time. In so doing, the citation index can gradually serve more and more and larger and larger humanities communities.

Creating the humanities citation index requires not only a growth strategy, but first and foremost a research infrastructure which provides for the right affordances to build the index as a collaborative, distributed and open effort. We propose its design in what follows.

4.5 Metadata ecosystem and requirements

In the research infrastructure needed to build an A&H citation index, libraries play a key role not only as holders of digitised collections but also as potential providers of data that can greatly support the citation extraction process. In fact, library catalogues constitute highly valuable knowledge bases of bibliographic information that can be exploited when doing citation mining, and especially citation matching.

We identify a set of key technical requirements that need to be met if library catalogue metadata are to be seamlessly integrated into the HuCI infrastructure. These requirements are:

1. Ability to handle the heterogeneity of metadata formats;
2. Provision of unique persistent identifiers;
3. Machine-aided creation, delivery and exchange of metadata;
4. Fine-grained/granular metadata descriptions;
5. Open licensing of metadata.

Metadata formats. From the point of view of citation mining pipelines and processes, there is a need to have metadata expressed in concise and “easy-to-process” formats. Such concerns become even more relevant when the metadata processing happens at a large scale, as the needs arise for optimising processing time and for efficient data storage. For example, in the context of previous work carried out by some of the authors [32,33], the Central Institute for the Union Catalogue of Italian Libraries and for Bibliographic Information (ICCU) has created a dump of 15 million records by transforming its data from MARC to a JSON-based representation, so as to facilitate their use in the project’s citation mining pipeline. Along similar lines, Bergamin and Bacchi [43] have successfully tested a workflow for mapping ICCU’s

UNIMARC data onto Wikibase Data Model, which would allow for using Wikibase as an environment to manage and edit bibliographic data, as well as exposing such data in an easier to process format.

MARC, however, is only one of the many formats that characterise the landscape of library metadata, where a plethora of old and new formats co-exist [44]. This situation makes it seem rather unlikely that libraries will converge to a common and widely adopted metadata format in the near future. As a result, a key requirement of the HuCI infrastructure is the ability to handle this heterogeneity of bibliographic metadata formats, achievable by developing code modules that read these formats and map them onto a common one.

Provision of unique identifiers. In addition to the granularity of descriptions, the provision of unique, persistent identifiers to identify bibliographic resources is another key requirement for metadata that are meant to support citation mining processes. Ideally, any primary or secondary source of which we are interested in tracking the citations ought to be identifiable by means of a unique, persistent identifier (e.g. a resolvable URI). Naturally, what is considered as a primary source varies from domain to domain: archival documents in History, various types of texts in Classics (e.g. canonical, papyri, inscriptions), manuscripts in Medieval Literature Studies, inscriptions and papyri in Egyptology, and so forth. Once these identifiers are in place, it is possible to use them to link ‘disambiguated’ citations. However, it cannot be the task of a single project to mint and provide these identifiers. This process should be happening in each discipline—and it has already been happening over the past years, e.g. in Classics [45]—but it can be fostered and accelerated by large-scale initiatives involving libraries and cultural heritage institutions, such as the European Open Science Cloud [46].

Machine-aided creation, delivery and exchange of metadata. There is an urgent need to take humans *out* of the loop insofar as access to and exchange of library metadata are concerned. Libraries—and especially aggregators of library metadata (e.g. national aggregators, library consortia, etc.)—ought to provide, at the very least, regular data dumps of their bibliographic metadata so as to facilitate their consumption and further reuse. Data dumps, however, being frozen snapshots of a dataset, raise the issue of synchronisation between the data at the source and the copy of the data used by other systems and processes. A partial solution to this problem is to provide streams of data (e.g. via APIs) in addition to regular dumps.

Granularity of bibliographic descriptions. The granularity of bibliographic descriptions is an apt example of gaps currently existing between the needs and requirements of citation mining projects, on the one hand, and the cataloguing practices currently adopted by the majority of libraries, on the other hand. Types of publications where granularity matters the most are journal articles, book chapters and indi-

vidual essays within collective volumes. In fact, while the citation unit of such publications is often the most granular (e.g. a given journal article, as opposed to the entire journal), cataloguing practices often do not reach that level of granularity in bibliographic descriptions.

Open licensing. Despite a declared willingness to share, often libraries and other cultural heritage institutions make available online data dumps that do not come with explicitly defined (open) licenses. They ought to be encouraged to always provide explicit license statements, as their absence hinders the reuse of shared data by others.

5 Research infrastructure

We propose to adopt a federated and distributed approach to design the research infrastructure required to create the Humanities Citation Index. Such an approach implies that the creation of citation data is delegated to a federation of cooperating institutions rather than being carried out by a single, central entity. The scenario we envisage (see Fig. 1) is having a network of GLAM institutions, each of them contributing citation data extracted from their digitised collections through a common open source software platform. These data will then be harvested, aggregated and consolidated to become the HuCI citation corpus, available to researchers both via search and exploration interfaces, as well as data dumps to be further analysed and visualised through external tools. In what follows we describe in more detail the proposed architecture, as well as the challenges related to its implementation.

5.1 Distributed and federated approach

A distributed and federated approach recognises the central role that libraries and other cultural heritage institutions could play with respect to the curation of their digitised collections. A modern notion of collection curation, we argue, ought to include the extraction of structured contents (e.g. citations) from digitised materials. This could take the form of manual verification, carried out by librarians, of automatically extracted information, as advocated by [47] for the specific case of citations.

At a technical level, a federated model has the advantage that it gives individual institutions a certain degree of freedom in deciding what can and should be made openly accessible—thus harvested from a federation of partner institutions—and what, on the contrary, should remain accessible only internally, within the boundaries of institutional access. We can call this access model “open by default and closed by necessity”. A typical case where it proves useful is the display of contextual information about citations (i.e. an excerpt of the text surrounding the citation), which can be complicated due

to copyright restrictions. A library holding digitised materials under copyright will want to give exclusively to its users full access to citation contexts from these publications, while still sharing with the wider community data extracted from these publications that do not fall under copyright, such as citation data. It is worth mentioning on this respect that citation data are just facts, and as such cannot be copyrighted. Thus, following the guidelines in [2] and the Initiative for Open Citations (I4OC),⁹ they should be released as public domain material using appropriate solutions, such as the CC0 waiver.

Moreover, a distributed approach makes sense not only for the collection of citation data but also for the hosting of the resulting citation corpus. In fact, collecting all the citation data from the whole scientific knowledge in one single, centralised repository—albeit feasible—would raise considerable issues in terms of maintenance and performances. For instance, consider that the data available in the OpenCitations Corpus include over 7.5 million bibliographic resources, mainly (80%) journal articles and their issues, volumes, and journals [48]. Their bibliographic metadata and their provenance information occupy more than 2 billion triples. Supposing that, roughly, 45 million new journal articles are published every year [49] and considering that 10% belong to the A&H, we can estimate 4.5 million new articles in the A&H every year. This amount of data can be stored in 230 gigabytes using the model adopted by OpenCitations—thus, we will need more than 5.7 terabytes for storing the metadata of all the A&H articles published in the “Web-era”, since 1994. While such amount of bytes can be even manageable in a big file system, these figures can drastically increase if we want to keep track of all the citation links among articles, if we start to ingest data coming from books (that are the primary publication object for the A&H which contains more citation links than any other scholarly medium), and if we consider publications that are older than 25 years.

The availability and cost of storage is something that will become more manageable in time. However, there is another, probably more pressing issue concerning the scenario mentioned above: having an infrastructure that guarantees efficient data querying is something very demanding when a large amount of data is actually available in a single and centralised database. As rough estimate considering the figures above, we would require 30 billions triples for handling the 4.5 million articles in A&H in the past 25 years. It is worth mentioning that this figure does not include books and all the older literature, nor any further extension of the kinds of data to store – for instance, the figures above are based on what is currently stored in the OpenCitations Corpus and, as such, do not account for abstracts or authors’ affiliations.

Therefore, it emerges that a centralised data storage solution is not feasible in the long-term.

5.2 Architecture

The HuCI infrastructure that we propose consists of four main layers, whose interplay is schematically depicted in Fig. 1:

1. The **data collection layer** consists of a network of content providers, who hold digitised materials (be they public or private) and contribute to the growth and coverage of the citation index by making openly available the citation data extracted from their holdings.
2. The **data federation layer** is conceived as a federation of decentralised citation databases based on RDF technologies where the HuCI citation data are actually stored.
3. The **service layer** provides HTTP APIs that allow for standardised access to HuCI data (e.g. via SPARQL endpoints and common REST Web APIs), external resources




Fig. 1 Four-layer architecture of the Humanities Citation Index. The solid arrows show that an item in one layer (ending node) uses the information provided by an item in another layer (starting node). The dashed line highlights existence of a federation mechanism between the items of the same layer

⁹ <https://i4oc.org>.

(e.g. those included in library or archive metadata) and services (e.g. author disambiguation).

4. Finally, the **application layer** is an ecosystem of tools and software components, plugged on HuCI's virtual triplestore (via the previous layer or by consuming directly the data from the data federation layer), that allow A&H researchers to discover and identify relevant literature for their research, and provides bibliometric insights into the citation data.

In what follows we discuss in greater detail each of these infrastructure components. In addition to such components, it is important that various providers of citation data are compliant as much as possible with the Principles of Open Scholarly Infrastructures (POSI) introduced in [50]. These principles are organised in three themes: Governance, Sustainability and Insurance. The latter theme specifies technological dimensions that should be guaranteed: **open source** (of all software required to run the infrastructure), **open data** (of all relevant data necessary to replicate it), **available data** (i.e. the availability of underlying data as periodic data dumps) and **patent non-assertion** (i.e. avoid using patents to prevent the community to replicate an infrastructure). If followed strictly, POSI should guarantee the long term sustainability of infrastructures that provide open scholarly data and open source software that can be used to build service new and innovative services. Several infrastructures (including OpenCitations, Crossref and DataCite) have run self-assessment exercises to measure their compliance with POSI, as introduced in the POSI website.¹⁰

5.2.1 Data collection

The first layer of HuCI's architecture is constituted by a network of GLAM institutions playing an active role in the production and curation of citation data. While each institution is responsible for the extraction of citation data from their digitised holdings, we envisage the development of a common open source platform that can ease the tasks of extracting citations from publications as well as the manual curation of such citations.

An example of software that could be deployed in the data collection layer is the Scholar Library (SL) platform¹¹ [32]. While it provides the typical functionalities of any digital library software (e.g. display of image and OCR), SL integrates specific components that perform the extraction of bibliographic references from digitised publications, and their disambiguation against bibliographic databases. In particular, it includes two components for the enrichment of publications with citation data: a machine learning-based

citation extractor as well as a component to match bibliographic references against the unified catalogue of Italian libraries [42].

The SL was designed to be *deployed locally* while staying *connected globally*: the local deployment ensures that digitised materials that cannot be shared openly remain private; APIs allow to harvest citation data from each local instance of SL and to connect them into a global citation index. As such, this platform could be deployed by partner institutions to facilitate the extraction and sharing of open citation data from their digitised holdings.

5.2.2 Data federation

As a long-term solution to devising a scalable infrastructure for the storage of A&H's citation data, we propose the HuCI virtual triple store, a federation of decentralised citation databases that can cooperate with each other by means of Web technologies, in particular RDF. Along the aforementioned lines, the interlinked databases of open citation data mentioned before have been recently released in order to address this aspect. The idea, in this aspect, is to organise existing and future open citations and scholarly metadata repositories (e.g. OpenCitations' datasets, Wikidata, OpenAIRE) as part of a bigger and interlinked graph of open repositories,¹² which would allow them to scale in terms of their infrastructure and the amount of data they need to handle. This can be implemented by means of appropriate Web and Semantic Web technologies, such as RDF triplestores, which natively are able to handle federated data storage. The use of such technologies is also crucial for enabling the development of a decentralised network of interoperable Linked Open Data (LOD), which are hosted in several places. In particular, such interoperability should be guaranteed by using the same data model for exposing the citation data involved, as introduced below. In essence, HuCI is a virtual database, since it must be implemented as a federation of repositories which provide access to their citation data via the HTTP protocol according to a particular shared data model, and which enable to expose the data in multiple formats (CSV, JSON, RDF-based formats) to foster maximum understandability for both humans and machines.

The feasibility of handling multiple and decentralised repositories of citation data effectively should be guaranteed by adopting a general metadata model in which the citation data will be described. If a different data model will be used by one of the repositories in the federated system, an explicit alignment to such general metadata model must be provided so as to make the federation possible.

¹⁰ <https://openscholarlyinfrastructure.org/posse>.

¹¹ <https://github.com/ScholarIndex/ScholarLibrary>.

¹² Something strongly supported by the 2017 report of the Confederation of Open Access Repositories (COAR), available at <https://www.coar-repositories.org/files/NGR-Final-Formatted-Report-cc.pdf>.

Among the possible candidate data models for describing citation data there is the OpenCitations Data Model (OCDM) [1]. OCDM is fully based on the Semantic Publishing and Referencing (SPAR) ontologies [51] and other standard vocabularies (FOAF, PROV, etc.) for the specification of additional information about agents and provenance data. The data model is implemented by means of the OpenCitations Ontology (OCO),¹³ which is not yet another bibliographic ontology, but rather simply a mechanism for grouping together existing complementary ontological entities from several other ontologies, for the purpose of providing descriptive metadata all in one place. As introduced in [52], the OCDM has already been adopted by several projects in the scholarly domain for organising bibliographic information such as the Venice Scholar Index¹⁴ [32], the Linked Open Citations Database (LOC-DB)¹⁵ [47] and the EXCITE Project¹⁶ [53].

5.2.3 Service layer

In addition to the two layers dedicated to data collection and federation, HuCI will comprise a layer of services (e.g. Web APIs) that will enable and regulate the flow of data between HuCI, its network of data providers, external providers of bibliographic metadata, and providers of services for the enrichment of citation data (both internal and external). In particular, we envisage three types of services:

1. services to harvest citation data from the network of data providers;
2. services to provide standardised access to external resources (e.g. archive and library catalogues);
3. services to enrich the aggregated citation data (e.g. interlinking, deduplication).

To the first type of services belong the APIs that will allow participating institutions to share the citation data extracted from their digitised collections. Citation data will be exposed by using the shared data model (such as the OpenCitations Data Model discussed above) and harvested via either SPARQL-based APIs or common Web REST APIs acting as a proxy to a SPARQL endpoint—that can be easily set up using software such as RAMOSE [54], BASIL [55], grlc [56], OBA [57], and SPARQL.anything [58]. These APIs could be available as part of HuCI or be offered by external providers. Provenance information, which includes the identification of the attribution, sources, activities and additional

change tracking data, is also attached to the related citation data in order to allow trackability and restorability of citation data due to some, even unpredictable, changes [59].

The second type of services aims to provide unified and standardised access to bibliographic metadata present in external resources, such as archive and library catalogues. These resources can be extremely valuable in various steps of the citation mining process (citation linking, author disambiguation), yet the heterogeneity of formats in which they are exposed hampers their reuse (see Sect. 4.5). These services will facilitate the access to external resources by defining a common API specification for data exchange, as well as a common data format towards which individual bibliographic formats can be mapped.

Finally, a third type of services will provide enrichment of citation data, especially through interlinking and deduplication. In fact, due to the federated nature of HuCI, it may happen that the same bibliographic entity and its citations are stored multiple times in different repositories. Thus, it is crucial to provide mechanisms and algorithms for dealing with deduplication appropriately, both for live access to data for *streaming* purposes, using a particular entry-point (e.g. a certain SPARQL endpoint), and to download full *dumps* of citation data available in different federated repositories.

The resolution of these conflicts could be handled by using persistent identification schemas (like DOI, Handle, ORCID or VIAF) for uniquely identifying various resources, or by applying disambiguation mechanisms based on entities' metadata. Of course, the more persistent identifiers are specified for a bibliographic entity, the easier its disambiguation will be and, consequently, the deduplication of bibliographic resources coming from different repositories.

A good example of integrating remote services into a common research infrastructure comes from the recent project Open Mining INFrastructure for TExt and Data (OpenMinTeD).¹⁷ Their API specification for processing Web services defines a protocol that allows remote NLP components to be seamlessly integrated into processing pipelines (see, e.g., [60]). Similarly, an API specification will need to be developed for external services that can be used to enrich HuCI's citation data.

5.2.4 Application layer

A crucial aspect of creating a citation index for the A&H concerns the development of user interfaces allowing researchers to explore and exploit citation data. In the technical infrastructure we propose, search and visualisation tools for the citation index will plug directly into HuCI's virtual triples store and will constitute its application layer. This layer will comprise user interfaces for search and visualisation, as well

¹³ <https://w3id.org/oc/ontology>.

¹⁴ <https://venicescholar.dhlab.epfl.ch>.

¹⁵ <https://locdb.bib.uni-mannheim.de>.

¹⁶ <http://excite.west.uni-koblenz.de>.

¹⁷ <http://openminted.eu>.

as software components that are meant to facilitate access to citation data stored in HuCI's virtual triples store via SPARQL API or via REST APIs built upon SPARQL endpoints.

Several tools have been developed to date to display, analyse and visualise citation data. They differ substantially with respect to the platform where they run (Web or desktop), their main purpose (analysis, visualisation, search), as well as the data sources for which they offer support (e.g. Web of Science, Scopus, PubMed, Microsoft Academics, OpenCitations, etc.). In particular, among these tools, there are VOSviewer [61], Sci2 [62], CiteSpace [63], Cytoscape [64], Bibliographic EXplorer (BEX) [65], CiteWiz [66], Docuipity [67], CRExplorer (CREx) [68], Science Citation Knowledge Extractor (SCKE) [69], Scholia [70], the Scholar Index (SI) [32], OSCAR [71], and LUCINDA.¹⁸

6 Conclusions

In this article, we have listed the main aspects that are necessary to devise the creation of a Humanities Citation Index (HuCI). We propose HuCI to be a decentralised and federated research infrastructure for gathering, sharing, elaborating, exposing bibliographic metadata and citation data of humanities publications that offers hooks for the development of further applications to keep track of the evolution of humanities research.

The technical guidelines we have provided for the creation of such an infrastructure follow current trends shared by the Open Science community around the globe. Several of the principles regarding data sharing we proposed are grounded in existing guidelines such as the FAIR (findability, accessibility, interoperability and reusability) data principles [72], which are considered a common and shared good practice in the field—where the word *data* in this context is an umbrella term including research data spreadsheets, software, workflows, slides and other research objects that accompany a *traditional* publication (e.g. a book, a journal article, a conference paper).

Several guidelines for enabling the creation of new open infrastructures—including their technological compliance, plans for their long-term sustainability and governance—have been proposed in the past five years, and have directly guided our work on HuCI. The Principles for Open Scholarly Infrastructures [73], the work done by the Confederation of Open Access Repositories (COAR) on best practices for implementing digital repositories [74,75] and other principles proposed by independent scholars such as the TRUST (transparency, responsibility, user focus, sustainability and technology) principles [76] have been extensively reused and

adapted to devise various component of the technical research infrastructure in HuCI. The very same principles characterise several national and international initiatives, such as the community workshop held in 2021 with the aim of shaping the main technical and organisational aspects for the creation of an open knowledge base of scholarly information for the Netherlands [77], and organisations created to help open infrastructures flourishing, such as the Global Sustainability Coalition for Open Science Services (SCOSS)¹⁹ and Invest in Open Infrastructures (IOI).²⁰

As part of our future work towards the creation of HuCI, we plan to conduct a survey among humanities scholars in order to elicit their views and desiderata with respect to the prospects and usefulness of such a citation index. This survey could be conducted in coordination with ongoing international activities on the topic of bibliographic data in the humanities, notably the DARIAH-EU Bibliographic Data Working Group.²¹ Nevertheless, given the striking similarities that citation indexes bear with thematic bibliographies (both printed and digital)—which are widely used by scholars across the humanities—it does seem plausible to postulate that such a citation index will meet the interests of many scholars.

Our hope is that the guidelines, principles and technological approaches described in this work can be an appropriate starting point for the implementation of HuCI, a fundamental tool for humanities research. The goals depicted by HuCI, and their technical implementation, are possible only if humanities scholars and institutions act together in a decentralised and coordinated fashion, by sharing efforts, resources and services towards a common objective, of which the suggestions in this article represent only the starting point.

Acknowledgements We thank the anonymous reviewers for their constructive feedback. The work of Matteo Romanello has been supported by the Swiss National Science Foundation under Grant Number PZ00P1_186033. The work of Silvio Peroni has been partially funded by the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 101017452.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copy-

¹⁸ <https://github.com/opencitations/lucinda>.

¹⁹ <https://scoss.org>.

²⁰ <https://investinopen.org>.

²¹ <https://www.dariah.eu/activities/working-groups/bibliographical-data-bibliodata>.

right holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Peroni, S., Shotton, D., Daquino, M.: The OpenCitations data model. Figshare (2019). <https://doi.org/10.6084/m9.figshare.3443876>
2. Peroni, S., Shotton, D.: Open citation: definition. Figshare (2018). <https://doi.org/10.6084/m9.figshare.6683855>
3. Stone, S.: Humanities scholars: information needs and uses. *J. Doc.* **38**(4), 292–313 (1982). <https://doi.org/10.1108/eb026734>
4. Ellis, D.: A behavioural model for information retrieval system design. *J. Inf. Sci.* **15**(4–5), 237–247 (1989). <https://doi.org/10.1177/016555158901500406>
5. Watson-Boone, R.: The information needs and habits of humanities scholars. *RQ* **34**(2), 203–215 (1994)
6. Wiberley, S.E., Jr.: Humanities literatures and their users. In: *Encyclopedia of Library and Information Sciences*, 3 edn, pp. 2197–2204 (2009). <http://hdl.handle.net/10027/7012>
7. Benardou, A., Constantopoulos, P., Dallas, C., Gavriliis, D.: Understanding the information requirements of arts and humanities scholarship. *Int. J. Digit. Curation* **5**(1), 18–33 (2010). <https://doi.org/10.2218/ijdc.v5i1.141>
8. Tibbo, H.R.: *Abstracting, Information Retrieval, and the Humanities: Providing Access to Historical Literature*. American Library Association, Chicago (1993)
9. Wiberley, S.E., Jones, W.G.: Patterns of information seeking in the humanities. *Coll. Res. Libr.* **50**(6), 638–645 (1989). https://doi.org/10.5860/crl_50_06_638
10. Bates, M.J.: The Getty end-user online searching project in the humanities: report no. 6: overview and conclusions. *Coll. Res. Libr.* **57**(6), 514–523 (1996)
11. Palmer, C., Tefteau, L., Pirmann, C.: Scholarly information practices in the online environment: themes from the literature and implications for library service development. Technical report (2009). <http://www.oclc.org/programs/publications/reports/2009-02.pdf>
12. Meho, L.I., Tibbo, H.R.: Modeling the information-seeking behavior of social scientists: Ellis's study revisited. *J. Am. Soc. Inf. Sci. Technol.* **54**(6), 570–587 (2003). <https://doi.org/10.1002/asi.10244>
13. Buchanan, G., Cunningham, S.J., Blandford, A., Rimmer, J., Warwick, C.: Information seeking by humanities scholars. In: Rauber, A., Christodoulakis, S., Tjoa, A.M. (eds.) *Research and Advanced Technology for Digital Libraries. Lecture Notes in Computer Science*, pp. 218–229. Springer, Berlin (2005). https://doi.org/10.1007/11551362_20
14. Martín-Martín, A., Thelwall, M., Orduna-Malea, E., Delgado López-Cózar, E.: Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations' COCI: a multidisciplinary comparison of coverage via citations. *Scientometrics* **126**(1), 871–906 (2021). <https://doi.org/10.1007/s11192-020-03690-4>
15. Visser, M., van Eck, N.J., Waltman, L.: Large-scale comparison of bibliographic data sources: Scopus, Web of Science, Dimensions, Crossref, and Microsoft Academic. *Quant. Sci. Stud.* **2**(1), 20–41 (2021). https://doi.org/10.1162/qss_a_00112
16. Harzing, A.-W., Alakangas, S.: Google Scholar, Scopus and the Web of Science: a longitudinal and cross-disciplinary comparison. *Scientometrics* **106**(2), 787–804 (2016). <https://doi.org/10.1007/s11192-015-1798-9>
17. Colavizza, G., Romanello, M.: Citation mining of humanities journals: the progress to date and the challenges ahead. *J. Eur. Period. Stud.* **4**(1), 36–53 (2019). <https://doi.org/10.21825/jeps.v4i1.10120>
18. Kulczycki, E., Engels, T.C.E., Pölönen, J., Bruun, K., Dušková, M., Guns, R., Nowotniak, R., Petr, M., Sivertsen, G., Isteni-Stari, A., Zuccala, A.: Publication patterns in the social sciences and humanities: evidence from eight European countries. *Scientometrics* **116**(1), 463–486 (2018). <https://doi.org/10.1007/s11192-018-2711-0>
19. Hicks, D.: The difficulty of achieving full coverage of international social science literature and the bibliometric consequences. *Scientometrics* **44**(2), 193–215 (1999)
20. Nederhof, A.J.: Bibliometric monitoring of research performance in the social sciences and the humanities: a review. *Scientometrics* **66**(1), 81–100 (2006). <https://doi.org/10.1007/s11192-006-0007-2>
21. Huang, M., Chang, Y.: Characteristics of research output in social sciences and humanities: from a research evaluation perspective. *J. Am. Soc. Inf. Sci. Technol.* **59**(11), 1819–1828 (2008). <https://doi.org/10.1002/asi.20885>
22. dos Santos, E.A., Peroni, S., Mucheroni, M.L.: Citing and referencing habits in medicine and social sciences journals in 2019. *J. Doc.* **77**(6), 1321–1342 (2021). <https://doi.org/10.1108/JD-08-2020-0144>
23. Freire, N., Robson, G., Howard, J.B., Manguinhas, H., Isaac, A.: Cultural heritage metadata aggregation using web technologies: IIIF, Sitemaps and Schema.org. *Int. J. Digit. Libr.* **21**(1), 19–30 (2020). <https://doi.org/10.1007/s00799-018-0259-5>
24. Ardanuy, J.: Sixty years of citation analysis studies in the humanities (1951–2010). *J. Am. Soc. Inf. Sci. Technol.* **64**(8), 1751–1755 (2013). <https://doi.org/10.1002/asi.22835>
25. Petr, M., Engels, T.C.E., Kulczycki, E., Dušková, M., Guns, R., Sieberová, M., Sivertsen, G.: Journal article publishing in the social sciences and humanities: a comparison of Web of Science coverage for five European countries. *PLoS ONE* **16**(4), 0249879 (2021). <https://doi.org/10.1371/journal.pone.0249879>
26. Hammarfelt, B.: Beyond coverage: toward a bibliometrics for the humanities. In: Ochsner, M., Hug, S.E., Daniel, H.-D. (eds.) *Research Assessment in the Humanities: Towards Criteria and Procedures*, pp. 115–131. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-29016-4_10
27. Heinzkill, R.: Characteristics of references in selected scholarly English literary journals. *Libr. Q.* **50**(3), 352–365 (1980)
28. Linmans, A.J.M.: Why with bibliometrics the humanities does not need to be the weakest link: indicators for research evaluation based on citations, library holdings, and productivity measures. *Scientometrics* **83**(2), 337–354 (2009). <https://doi.org/10.1007/s11192-009-0088-9>
29. Sula, C.A., Miller, M.: Citations, contexts, and humanistic discourse: toward automatic extraction and classification. *Lit. Linguist. Comput.* **29**(3), 452–464 (2014). <https://doi.org/10.1093/llc/fqu019>
30. Chen, S.-C.: Exploring the use of electronic resources by humanities scholars during the research process. *Electron. Libr.* **37**(2), 240–254 (2019). <https://doi.org/10.1108/EL-08-2018-0170>
31. Knievel, J., Kellsey, C.: Citation analysis for collection development: a comparative study of eight humanities fields. *Libr. Q.* **75**(2), 142–168 (2005). <https://doi.org/10.1086/431331>
32. Colavizza, G., Romanello, M., Babetto, M., Barbay, V., Bolli, L., Ferronato, S., Kaplan, F.: *Linked Books: Towards A Collaborative Citation Index for the Arts and Humanities*, pp. 178–181. *Red de Humanidades Digitales, A. C.*, Mexico City (2018)
33. Colavizza, G., Romanello, M.: Citation mining of humanities journals: the progress to date and the challenges ahead. *J. Eur. Period. Stud.* **4**, 36–53 (2019)

34. Franssen, T., Wouters, P.: Science and its significant other: representing the humanities in bibliometric scholarship. *J. Am. Soc. Inf. Sci.* **70**(10), 1124–1137 (2019). <https://doi.org/10.1002/asi.24206>
35. Waltman, L.: A review of the literature on citation impact indicators. *J. Informetr.* **10**(2), 365–391 (2016). <https://doi.org/10.1016/j.joi.2016.02.007>
36. Hicks, D., Wouters, P., Waltman, L., de Rijcke, S., Rafols, I.: Bibliometrics: the Leiden Manifesto for research metrics. *Nature* **520**(7548), 429–431 (2015). <https://doi.org/10.1038/520429a>
37. Scheidel, W.: Continuity and change in classical scholarship. *Ancient Soc.* **28**, 265–289 (1997). <https://doi.org/10.2143/AS.28.0.630079>
38. Federation Of Finnish Learned Societies, Information, T.C.F.P., Publishing, T.F.A.F.S., Universities Norway, European Network for Research Evaluation in the Social Sciences and the Humanities: Helsinki initiative on multilingualism in scholarly communication. Figshare (2019). <https://doi.org/10.6084/m9.figshare.7887059>
39. IFLA: Functional requirements for bibliographic records (1997). https://www.ifla.org/wp-content/uploads/2019/05/assets/cataloguing/frbr/frbr_2008.pdf
40. Hellqvist, B.: Referencing in the humanities and its implications for citation analysis. *J. Am. Soc. Inf. Sci. Technol.* **61**(2), 310–318 (2010). <https://doi.org/10.1002/asi.21256>
41. Grafton, A.: *The Footnote: A Curious History*. Harvard University Press, Cambridge, MA (1999). OCLC: 1264778206
42. Colavizza, G., Romanello, M., Kaplan, F.: The references of references: a method to enrich humanities library catalogs with citation data. *Int. J. Digit. Libr.* **19**(2–3), 151–161 (2018). <https://doi.org/10.1007/s00799-017-0210-1>
43. Bergamin, G., Bacchi, C.: New ways of creating and sharing bibliographic information: an experiment of using the Wikibase Data Model for UNIMARC data. *JLIS.it* **9**(3), 35–74 (2018). <https://doi.org/10.4403/jlis.it-12458>
44. Tennant, R.: A bibliographic metadata infrastructure for the twenty-first century. *Libr. Hi Tech* **22**(2), 175–181 (2004). <https://doi.org/10.1108/07378830410524602>
45. Romanello, M., Pasin, M.: Using linked open data to bootstrap a knowledge base of classical texts. In: Adamou, A., Daga, E., Isaksen, L. (eds.) *Proceedings of the Second Workshop on Humanities in the Semantic Web (WHiSe)*. CEUR Workshop Proceedings, Aachen, p. 12 (2017). <http://ceur-ws.org/Vol-2014/paper-01.pdf>
46. Hellström, M., Heughebaert, A., Kotarski, R., Manghi, P., Matthews, B., Ritz, R., Conrad, A.S., Weigel, T., Wittenburg, P., Valle, M.: Second draft Persistent Identifier (PID) policy for the European Open Science Cloud (EOSC) (2020). <https://doi.org/10.5281/zenodo.3780423>
47. Lauscher, A., Eckert, K., Galke, L., Scherp, A., Rizvi, S.T.R., Ahmed, S., Dengel, A., Zumstein, P., Klein, A.: Linked open citation database: enabling libraries to contribute to an open and interconnected citation graph. In: *Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries—JCDL '18*, pp. 109–118. ACM Press, Fort Worth (2018). <https://doi.org/10.1145/3197026.3197050>
48. Peroni, S., Shotton, D., Vitali, F.: One year of the OpenCitations Corpus. In: d'Amato, C., Fernandez, M., Tamma, V., Lecue, F., Cudré-Mauroux, P., Sequeda, J., Lange, C., Heflin, J. (eds.) *The Semantic Web—ISWC 2017*. Lecture Notes in Computer Science, vol. 10588, pp. 184–192. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68204-4_19
49. van Harmelen, F.: *The end of the scientific paper as we know it (or not...)*. Wien, Austria (2017). <https://www.slideshare.net/Frank.van.Harmelen/the-end-of-the-scientific-paper-as-we-know-it-or-not>
50. Bilder, G., Lin, J., Neylon, C.: Principles for open scholarly infrastructures. Figshare (2015). <https://doi.org/10.6084/m9.figshare.1314859>
51. Peroni, S., Shotton, D.: The SPAR ontologies. In: Rutkowski, L., Scherer, R., Korytkowski, M., Pedrycz, W., Tadeusiewicz, R., Zurada, J.M. (eds.) *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. 10842, pp. 119–136. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00668-6_8
52. Daquino, M., Peroni, S., Shotton, D., Colavizza, G., Ghavimi, B., Lauscher, A., Mayr, P., Romanello, M., Zumstein, P.: The OpenCitations data model. In: Pan, J.Z., Tamma, V., d'Amato, C., Janowicz, K., Fu, B., Polleres, A., Seneviratne, O., Kagal, L. (eds.) *The Semantic Web—ISWC 2020*. Lecture Notes in Computer Science, vol. 12507, pp. 447–463. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-62466-8_28
53. Hosseini, A., Ghavimi, B., Boukhers, Z., Mayr, P.: EXCITE—a toolchain to extract, match and publish open literature references, Urbana-Champaign, Illinois, USA (2019). <https://philippmayr.github.io/papers/JCDL2019-EXCITE-demo.pdf>
54. Daquino, M., Heibi, I., Peroni, S., Shotton, D.: Creating Restful APIs over SPARQL endpoints with RAMOSE. In: *Semantic Web (2020)*
55. Daga, E., Panziera, L., Pedrinaci, C.: BASIL: A cloud platform for sharing and reusing SPARQL queries as Web APIs. In: Vilalta, S., Pan, J.Z., Dragoni, M. (eds.) *Proceedings of the ISWC 2015 Posters & Demonstrations Track Co-located with the 14th International Semantic Web Conference (ISWC-2015)*, vol. 1486, p. 4. CEUR-WS, Bethlehem (2015). http://ceur-ws.org/Vol-1486/paper_41.pdf
56. Meroño-Peñuela, A., Hoekstra, R.: grlc makes GitHub taste like linked data APIs. In: Sack, H., Rizzo, G., Steinmetz, N., Mladenčić, D., Auer, S., Lange, C. (eds.) *The Semantic Web*. Lecture Notes in Computer Science, vol. 9989, pp. 342–353. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47602-5_48
57. Garjijo, D., Osorio, M.: OBA: an ontology-based framework for creating REST APIs for knowledge graphs. In: Pan, J.Z., Tamma, V., d'Amato, C., Janowicz, K., Fu, B., Polleres, A., Seneviratne, O., Kagal, L. (eds.) *The Semantic Web—ISWC 2020*. Lecture Notes in Computer Science, vol. 12507, pp. 48–64. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-62466-8_4
58. Daga, E., Asprino, L., Mulholland, P., Gangemi, A.: Facade-X: an opinionated approach to SPARQL anything. [arXiv:2106.02361](https://arxiv.org/abs/2106.02361) [cs]
59. Peroni, S., Shotton, D., Vitali, F.: A document-inspired way for tracking changes of RDF data. In: *Proceedings of the 1st Workshop on Detection, Representation and Management of Concept Drift in Linked Open Data*. CEUR Workshop Proceedings, Aachen, p. 8 (2016). http://ceur-ws.org/Vol-1799/Drift-a-LOD2016_paper_4.pdf
60. Ba, M., Bossy, R.: Interoperability of Corpus Processing Workflow Engines: The Case of AlvisNLP/ML in OpenMinTeD. *Zenodo* (2016). <https://doi.org/10.5281/zenodo.200370>
61. van Eck, N.J., Waltman, L.: Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics* **84**(2), 523–538 (2010). <https://doi.org/10.1007/s11192-009-0146-3>
62. Team, S.: *Sci2 Tool: A Tool for Science of Science Research and Practice*. Indiana University and SciTech Strategies (2009). <https://sci2.cns.iu.edu/user/index.php>
63. Chen, C.: Searching for intellectual turning points: progressive knowledge domain visualization. *Proc. Natl. Acad. Sci.* **101**(Supplement 1), 5303–5310 (2004). <https://doi.org/10.1073/pnas.0307513100>
64. Shannon, P.: Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res.* **13**(11), 2498–2504 (2003). <https://doi.org/10.1101/gr.1239303>
65. Di Iorio, A., Giannella, R., Poggi, F., Peroni, S., Vitali, F.: Exploring scholarly papers through citations. In: *Proceedings of the*

- 2015 ACM Symposium on Document Engineering. DocEng '15, pp. 107–116. Association for Computing Machinery, New York (2015). <https://doi.org/10.1145/2682571.2797065>
66. Elmqvist, N., Tsigas, P.: CiteWiz: a tool for the visualization of scientific citation networks. *Inf. Vis.* **6**(3), 215–232 (2007). <https://doi.org/10.1057/palgrave.ivs.9500156>
 67. Poggi, F., Ciancarini, P., Iorio, A.D., Peroni, S., Vitali, F.: Exploiting Coordinated Views for Scholarly Reading and Analysis, pp. 113–124 (2019). <https://doi.org/10.18293/DMSVIVA2019-021>
 68. Thor, A., Marx, W., Leydesdorff, L., Bornmann, L.: Introducing CitedReferencesExplorer (CRExplorer): a program for reference publication year spectroscopy with cited references standardization. *J. Informetr.* **10**(2), 503–515 (2016). <https://doi.org/10.1016/j.joi.2016.02.005>
 69. Lent, H., Hahn-Powell, G., Haug-Baltzell, A., Davey, S., Surdeanu, M., Lyons, E.: Science citation knowledge extractor. *Front. Res. Metr. Anal.* **3**, 35 (2018). <https://doi.org/10.3389/frma.2018.00035>
 70. Nielsen, F.A., Mietchen, D., Willighagen, E.: Scholia, scientometrics and wikidata. In: Blomqvist, E., Hose, K., Paulheim, H., Ławrynowicz, A., Ciravegna, F., Hartig, O. (eds.) *The Semantic Web: ESWC 2017 Satellite Events. Lecture Notes in Computer Science*, pp. 237–259. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70407-4_36
 71. Heibi, I., Peroni, S., Shotton, D.: Enabling text search on SPARQL endpoints through OSCAR. *Data Sci.* (2019). <https://doi.org/10.3233/DS-190016>
 72. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L.B., Bourne, P.E., Bouwman, J., Brookes, A.J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C.T., Finkers, R., Gonzalez-Beltran, A., Gray, A.J.G., Groth, P., Goble, C., Grethe, J.S., Heringa, J., 't Hoen, P.A.C., Hooft, R., Kuhn, T., Kok, R., Kok, J., Lusher, S.J., Martone, M.E., Mons, A., Packer, A.L., Persson, B., Rocca-Serra, P., Roos, M., van Schaik, R., Sansone, S.-A., Schultes, E., Sengstag, T., Slater, T., Strawn, G., Swertz, M.A., Thompson, M., van der Lei, J., van Mulligen, E., Velterop, J., Waagmeester, A., Wittenburg, P., Wolstencroft, K., Zhao, J., Mons, B.: The FAIR Guiding Principles for scientific data management and stewardship. *Sci. Data* **3**, 160018 (2016). <https://doi.org/10.1038/sdata.2016.18>
 73. Bilder, G., Lin, J., Neylon, C.: *The Principles of Open Scholarly Infrastructure* (2020). <https://doi.org/10.24343/C34W2H>
 74. COAR WG Next Generation Repositories: Behaviours and Technical Recommendations of the COAR Next Generation Repositories Working Group. Recommendation, COAR (2017). <https://www.coar-repositories.org/files/NGR-Final-Formatted-Report-cc.pdf>
 75. Confederation of Open Access Repositories: COAR community framework for best practices in repositories. Technical Report Version 1, Zenodo (2020). <https://doi.org/10.5281/zenodo.4110829>
 76. Lin, D., Crabtree, J., Dillo, I., Downs, R.R., Edmunds, R., Giaretta, D., De Giusti, M., L'Hours, H., Hugo, W., Jenkyns, R., Khodiyar, V., Martone, M.E., Mokrane, M., Navale, V., Peters, J., Sierman, B., Sokolova, D.V., Stockhause, M., Westbrook, J.: *The TRUST Principles for digital repositories*. *Sci. Data* **7**(1), 144 (2020). <https://doi.org/10.1038/s41597-020-0486-7>
 77. Neylon, C., Bijsterbosch, M., Dunning, A., Kramer, B., De Rijcke, S., Tatum, C., Waltman, L.: *An open knowledge base for the Netherlands: report of a community workshop*. Technical report, Zenodo (2021). <https://doi.org/10.5281/ZENODO.4893803>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.