A survey of Web technology for metadata aggregation in cultural heritage

Nuno Freire [a,*], Antoine Isaac [b], Glen Robson [c], John Brooks Howard [d] and
Hugo Manguinhas [b]

[a] *INESC-ID, Rua Alves Redol 9, 1000-029 Lisboa, Portugal*
E-mail: nuno.freire@tecnico.ulisboa.pt
[b] *Europeana Foundation, Willem-Alexanderhof 5, 2509 LK The Hague, The Netherlands*
E-mails: antoine.isaac@europeana.eu, hugo.manguinhas@europeana.eu
[c] *National Library of Wales, Penglais Rd, Aberystwyth SY23 3BU, United Kingdom*
E-mail: glen.robson@llgc.org.uk
[d] *University College Dublin, Stillorgan Rd, Belfield, Dublin 4, Ireland*
E-mail: john.b.howard@ucd.ie

Abstract. In the World Wide Web, a very large number of resources are made available through digital libraries. The existence of many individual digital libraries, maintained by different organizations, brings challenges to the discoverability and usage of these resources by potential users. A widely-used approach is metadata aggregation, where a central organization takes the role of facilitating the discoverability and use of the resources, by collecting their associated metadata. The central organization has the possibility to further promote the usage of the resources by means that cannot be efficiently undertaken by each digital library in isolation. This paper focuses in the domain of cultural heritage, where OAI-PMH has been the embraced solution, since discovery of resources was only feasible if based on metadata instead of full-text. However, the technological landscape has changed. Nowadays, with the technological improvements accomplished by network communications, computational capacity, and Internet search engines, the motivation for adopting OAI-PMH is not as clear as it used to be. In this paper, we present the results of our analysis of available potential technologies, using as application context the Europeana Network and its requirements for metadata aggregation. We cover the following technologies: IIIF (International Image Interoperability Framework); Webmention; Linked Data Notifications; WebSub; Sitemaps; ResourceSync; Open Publication Distribution System (OPDS); Linked Data Platform; and Schema.org.

Keywords: Metadata, cultural heritage, linked data, Web technology, standards

1. Introduction

In the World Wide Web, a very large number of resources is made available through digital libraries. The existence of many individual digital libraries, maintained by different organizations, brings challenges to the discoverability and usage of the resources by potential interested users.

An often-used approach is metadata aggregation, where a central organization takes the role of facilitating the discovery and use of the resources by collecting their associated metadata. Based on these aggregated datasets of metadata, the central organization (often called aggregator) can further promote the usage of the resources by means that cannot be efficiently undertaken by each digital library in

---

*Corresponding author. E-mail: nuno.freire@tecnico.ulisboa.pt; Tel.: +351.213100300; Fax: +351.213145843.

isolation. This scenario is widely applied in the domain of cultural heritage, where the number of organizations with their own digital libraries is very large. In Europe, Europeana has the role of facilitating the usage of cultural heritage resources from and about Europe, and although many European cultural heritage organizations do not yet have a presence in Europeana, it already holds metadata of resources originating from more than 3,500 providers (source: http://statistics.europeana.eu/europeana [consulted on 4th of January 2017]).

This domain is also characterized by users that often have very specific information needs, which cannot be easily fulfilled by the Internet search engines. The retrieval of resources based on metadata, in combination with the hypertext documents of the World Wide Web, has been a challenge that the search engines have not yet been able to provide an effective solution for, therefore the retrieval of cultural heritage resources via search engines is ineffective.

The technological approach to metadata aggregation has been mostly based on the OAI-PMH protocol, a technology initially designed in 1999. OAI-PMH was meant to address shortcomings in scholarly communication by providing a technical interoperability solution for discovery of e-prints, via metadata aggregation. The cultural heritage domain embraced the solution offered by OAI-PMH, however, the technological landscape around our domain has changed. Nowadays, cultural heritage organizations are increasingly applying technologies designed for the wider interoperability on the World Wide Web. Particularly relevant for our work are those related with the social web, the web of data, Internet search engine optimization, and the IIIF (International Image Interoperability Framework).

In this paper, we present the results of our work in surveying available web technology for applicability in metadata aggregation in cultural heritage. This work is part of our aim to rethink the technological approach for metadata aggregation, with the goal of finding a solution to make the continuous operation of aggregations networks more efficient and to lower the technical barriers for data providers to share their resources.

Our work is guided by the study of the existing aggregation network of Europeana, from where we identify the requirements for metadata aggregation. Europeana provides access to digitised cultural resources from a wide range of cultural heritage institutions across Europe, mostly including libraries, museums, archives and galleries. It seeks to enable users to search and access knowledge in all the languages of Europe. This is done either directly, via its web portals, or indirectly, via third-party applications built on top of its data services (search APIs and Linked Open Data).

The Europeana service is based on the aggregation and exploitation of (meta)data about the digitized objects from very different contexts. To provide a seamless, efficient services on top of such aggregation, it must solve hard data integration issues. To address these, Europeana has developed infrastructures and workflows for aggregating, ingesting, indexing, normalising, and publishing data.

This paper makes the following scientific contributions to the digital libraries community:

- An analysis of requirements for metadata aggregation based on a large network of data providers – the Europeana Network.
- A functional analysis for innovative use of state of the art technologies.
- A real-world application experience of open standards, thus contributing for their future improvement.

The paper will describe, in Section 2, the technological approach to metadata aggregation most prevalent in cultural heritage. Specific requirements, which guided our technological survey, are presented in Section 3. The Web technologies that were analyzed are presented in Section 4, and Section 5 concludes.

2. Metadata aggregation in cultural heritage – past and present

In the cultural heritage domain, the technological approach to metadata aggregation has been mostly based on the OAI-PMH protocol, a technology initially designed in 1999 [9]. OAI-PMH was originally meant to address shortcomings in scholarly communication by providing a technical interoperability solution for discovery of e-prints, via metadata aggregation.

The cultural heritage domain embraced the solution offered by OAI-PMH, since discovery of resources was only feasible if based on metadata instead of full-text [18]. In Europe, OAI-PMH had one of its largest, and earliest, applications in The European Library [19], which aggregated digital collections and bibliographic catalogues from 48 national libraries. It was also the technological solution adopted by Europeana since its start, to aggregate metadata from its network of data providers and intermediary aggregators [12].

However, the technological landscape around our domain has changed. Nowadays, with the technological improvements accomplished by network communications, computational capacity, and Internet search engines, the discovery of resources, such as e-prints, is largely based on full-text processing, thus the newer technical advances, such as ResourceSync [10], are less focused on metadata. Within the cultural heritage domain metadata-based discovery remains the most widely adopted approach since a lot of material is not available as full-text. The adoption of OAI-PMH for this purpose is not as clear as it used to be, however. OAI-PMH was designed before the key founding concepts of the Web of Data [2]. By being centered on the concept of repository, instead of centering on the resources, the protocol is often misunderstood and its implementations fail, or are deployed with flaws that undermine its reliability [18]. Another important factor is that OAI-PMH predates REST [13]. Thus, it does not follow the REST principles, further bringing resistance and difficulties in its comprehension and implementation by developers in cultural heritage organizations.

An additional aspect relevant for our work, is that nowadays, cultural heritage organizations are increasingly applying technologies designed for wider interoperability on the World Wide Web. Particularly relevant are those related with Internet search engine optimization and the International Image Interoperability Framework [16]. Regardless of the metadata aggregation process for Europeana, cultural heritage institutions are already interested in developing their systems' capabilities in these areas. By exploring these technologies, the participation in Europeana of these institutions may become much less demanding and possibly even transparent.

The cultural heritage domain has some specific characteristics, which have heavily influenced how metadata aggregation has been conducted in the past. We consider the following to be the most influential:

- Several sub domains compose the cultural heritage domain: Libraries, Archives, and Museums (the term LAM is often used to refer to the three sub domains).
- Interoperability of systems and data is scarce across sub-domains, but it is common within each sub-domain, both at the national and the international level.
- Each sub-domain 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. A long-time standardization tradition has existed in libraries, while this practice is more recent in archives and museums.
- Several of the adopted standards tend to be flexible towards data structure. Standards based on relational data models, for example, are rare in cultural heritage, while XML-based data models are common.

- Organizations typically have limited budgets to devote to information and communication technologies, thus the speed and extent of innovation and adoption of new technologies is slow.

In this environment, a common practice has been to aggregate metadata, under an agreed data model that allows the data heterogeneity between organizations and countries to be dealt with in a sustainable way. These data models typically address two main requirements:

- Retaining the semantics of the original data from the source providers
- Supporting the information needs of the services provided by the aggregator.

These two requirements are typically addressed in a way that keeps the model complexity low, with the intention of simplifying the understanding of the model by all kinds of providers, and to allow for a low barrier of implementation of data conversion solutions, by both providers and aggregators.

Another relevant aspect of metadata aggregation is the sharing of the sets of metadata from the providing organizations to the aggregator. The metadata is transferred to the aggregator, but it continues to evolve at the data provider, thus the aggregator needs to periodically update its copy of the data. In this case, the needs for data sharing can be described as a cross-organizational data synchronization problem.

In the cultural heritage domain, OAI-PMH is the most well-established solution to address the data synchronization problem. Since OAI-PMH is not restrictive in terms of the data model to be used, it allows the sharing of the metadata per the adopted data model of each aggregation case. The only restriction imposed by OAI-PMH is that the metadata must be represented in XML.

In the case of Europeana, the technological solutions for the aspects of data modelling and data synchronization have evolved at different rhythms. While for data modelling the Europeana Data Model (EDM) [1] has always been under continuous improvement, the solution for data synchronization, based on OAI-PMH, has not been reassessed since its early adoption.

Another important aspect of EDM, is that it does not impose any constraint in the choice of Web technologies for data synchronization. This comes from EDM following the principles of the Web of Data, and that it can be serialized in XML and in RDF formats. This aspect gives the Europeana network much choice for technological innovation of the aggregation network.

3. Requirements of the Europeana Network

Data aggregation is a general information systems problem, for which computer scientists have provided many possible solutions. The type of solution applicable to each case is greatly influenced by the requirements of the application scenario.

The Europeana Network is a network of data providing CHIs. When addressing aggregation across organizations, the technological capacity of the participating organizations is a key determinant of the solution to be applied. We define the requirements for the solution by considering the characteristics of the cultural heritage domain along with some particularities of the metadata aggregation carried out in the Europeana Network by data providers and aggregators, and the legacy of the current established aggregation practices.

This section presents the requirements separated between the two sub-problems of data aggregation: the synchronization of data sources, and data modeling/representation.

3.1. Synchronization of data sources

The type of solution for synchronization of data sources across organizations is greatly influenced by the requirements for data consistency guarantees, and synchronization latency. For the Europeana Network the solution must allow an aggregator to collect structured metadata about the digital resources that a CHI (the provider) wants to make available in Europeana. A solution should address the following requirements:

- The set of resources for aggregation is specified by the provider, and may comprehend all the resources of a digital library, or just a subset.
- The set of aggregated resources may evolve over time; therefore, the synchronization process must provide efficient mechanisms for incremental aggregation.
- The synchronization process between the provider and Europeana must be automatic and efficient, in terms of computation and network communication.
- The synchronization mechanism must be scalable to the level of the largest datasets nowadays available in Europeana, which are in the range of 2–5 million resources.
- A solution should be simple to adopt by data providers. One of the following aspects would make a solution simple to adopt:

  - It is based on technologies already in use by data providers;
  - It has very simple technical requirements for implementation;
  - Open source and free tools exist for deploying the solution.

- The solution can be more technologically challenging on the aggregators side than on the data providers', since the aggregators are often better prepared to address more complex technical implementation issues of information systems.

In the context of the above requirements, Section 4 presents the Web technologies that we identified as possible solutions for data synchronization.

3.2. Data modeling and representation

In the current Europeana aggregation network, EDM is the technology that supports data sharing efforts in the aspect of data modeling and representation. It is a solution that allows Europeana to become 'a big aggregation of digital representations of culture artefacts together with rich contextualization data and embedded in a linked Open Data architecture' [8].

EDM also has a key role in many other parts of the Europeana Network. EDM has been a collaborative effort from the very start, involving representatives from all the domains represented in Europeana: libraries, museums, archives and galleries. It supports several of the core processes of the Europeana's operations, and contributes to the access layer of the Europeana Platform, where it supports the data reuse by third parties [4]. EDM's influence and usage also reaches beyond the Europeana Network, with notable cases such as the Digital Public Library of America that defined its metadata application profile based on EDM [5].

Although reducing the data conversion effort required within the Europeana aggregation infrastructure is a very relevant aspect, in our work we considered that any innovative mechanisms for data modeling/representation to be feasible for application in Europeana, should not impact the other areas where EDM is used. We therefore consider that any new technological solution should address the following requirements:

- It must have the capacity to represent the required information for the minimal requirements of EDM and high-quality cultural heritage data;
- It should be flexible, especially not committed to the vision of only one of the cultural heritage domains Europeana is serving, and ideally offering an easy implementation/learning curve (e.g., allowing Dublin Core-level expression of metadata);
- It should show signs of significant adoption and/or interest.

Given the above requirements, searching for innovative technologies in this area has proven to be very hard in the course of our work. Only one viable solution was identified – Schema.org, which is described in the Section 4.9.

4. Web technologies for metadata aggregation

Most of the technologies described in this section were designed for fulfilling the needs of general use cases, and are applicable across several domains. Some of these can completely fulfil the requirements of metadata aggregation, while others only do so partially, and need to be combined with other technologies. Not all technologies have been explored, in our work, to the same level of detail, but, in this section, we describe all those that we have identified as being applicable.

4.1. International image interoperability framework

The International Image Interoperability Framework, commonly known as IIIF, is a family of specifications that were conceived to facilitate systematic reuse of image resources in digital image repositories maintained by cultural heritage organizations. It specifies several HTTP based web services [16] covering access to images, the presentation and structure of complex digital objects, composed of one or more images, and searching within their content.

IIIF strength resides in the presentation possibilities it provides for end-users. From the perspective of data acquisition, however, none of the IIIF APIs was specifically designed to support metadata aggregation. Nevertheless, within the output given by the IIIF APIs, there may exist enough information to allow HTTP robots to crawl IIIF endpoints and harvest the links to the digital resources and associated metadata.

To study the feasibility of data acquisition via IIIF, several experiments and case studies have been undertaken, and are currently in progress. The early experiments revealed that IIIF contains all the necessary elements for automatic harvesting of metadata. Some of these elements are, however, not of mandatory implementation, thus they will not be available in many IIIF endpoints. The following elements of IIIF APIs must be provided by data providers, to enable Europeana to harvest:

- Structured metadata: the typical metadata available in the output of IIIF is intended for end-user presentation, thus it is unable to fulfil the requirements of ingestion in Europeana. This limitation may however be overcome by using the optional links (i.e. seeAlso) to structured metadata, as specified in IIIF. These enable crawlers to harvest metadata in any format provided, such as EDM, Dublin Core, etc.
- IIIF Collection indicating the resources for Europeana: In IIIF, it is not required that the endpoint implements a mechanism to make publicly known all the digital objects that it makes available. However, such mechanism may be implemented, and, optionally, the IIIF provider may implement a IIIF Collection that lists the digital objects it holds, or just those intended for delivery to Europeana.

By making this collection known to Europeana, all the digital objects referenced in the collection can be crawled, and their metadata harvested by Europeana.

There is one piece of information that IIIF does not provide, which is the modification timestamp of the digital objects. This aspect has an impact in the efficiency of the harvesting process, but only becomes relevant in very large collections, with sizes in the hundreds of thousands of digital objects. In the typical size of the collections delivered to Europeana, within the thousands or tens of thousands, the loss in efficiency is not significant nowadays, due to high availability of bandwidth and computational capacity.

To overcome this issue of harvesting efficiency in large collections, other technologies may be used in conjunction with IIIF. Examples are Sitemaps, HTTP Headers, and notification protocols, such as Webmention and Linked Data Notifications, which we are also being evaluated in our work and are described in this document. This issue of harvesting efficiency has been brought to the attention of the IIIF community, and we are engaged in the discussions for achieving a standard mechanism, or recommendations, which will address it within the IIIF community.

The results so far indicate that data acquisition via IIIF is feasible, and presents little technological barriers for data providers that already have an IIIF solution in place for their own purposes. In the Europeana side, once a IIIF crawler tool is integrated with its aggregation management system, ingestion of IIIF data sources can be carried out under the same process of nowadays.

4.2. Sitemaps

Sitemaps [14] allow webmasters to inform search engines about pages on their sites that are available for crawling by search engine's robots. A Sitemap is an XML file that lists URLs of the pages within a website along with additional metadata about each URL (i.e., when it was last updated, how often it usually changes, and how important it is, relative to other URLs within the same site) so that search engines can more efficiently crawl the site. Sitemaps is a widely-adopted technology, supported by all major search engines. Many content management systems support Sitemaps out-of-the-box, and Sitemaps are simple enough to be manually built by webmasters when necessary.

Considering the application of Sitemaps in the context of Europeana, for data acquisition, it presents the following positive points:

- A simple technology with low barriers for implementation, even for small organizations.
- Already in use in several cultural heritage organizations, where it is applied for search engine optimization of their websites and digital libraries.
- It is extensible; thus, it can be adapted to Europeana specific requirements. For example, Google has Sitemap extensions for images and for videos, each one defining a set of metadata elements for its media type.

A Sitemap is an XML file, which is prepared per the Sitemap Protocol [14]. In digital libraries, Sitemaps typically contain all the links to the landing pages of the digital objects within the digital library.

These kinds of Sitemaps are widely used, thus already existing Sitemaps could be used by Europeana for metadata aggregation, using a WebCrawler such as those used by Internet search engines. Starting by following the links in a Sitemap, and processing structured data within HTML (e.g. microdata, Schema.org, linked data available by content negotiation), an Europeana Crawler may discover the digital cultural heritage objects, as well as metadata.

Besides its typical use for Internet crawlers, Sitemaps may also be deployed by Europeana and data providers in conjunction with other technologies, which would allow for simple ways of sharing data. For example, Sitemaps could be made available by data providers, in order to inform Europeana of the digital objects to be aggregated and when they are updated.

Sitemaps, present two clear benefits: a very low technological barrier, and data providing organizations often have in-house knowledge about XML and/or Sitemaps. Sitemaps are a key technology applied for Internet search engine optimization, thus it is already in use within data providers' websites and digital libraries for making their resources discoverable in Internet search engines. Providing metadata to Europeana by using Sitemaps would substantially reduce the implementation effort needed by data providers.

4.3. ResourceSync

ResourceSync [10] is a NISO standard that enables third-party systems to remain synchronized with a data provider's evolving digital objects, supporting both metadata and content. ResourceSync is based on the Sitemaps protocol and introduces extensions that enable its functionality for accurately and efficiently synchronizing the content of digital objects. Additionally, to Sitemap's capabilities, it allows data sources to:

- specify groups of resources, instead of each one individually.
- specify alternative ways to download the resources, as for example, as a bundle in a zip file.
- specify what has changed at a time.
- specify alternative ways to download just a set of changes.
- link resources to metadata that describes the resources.
- link to older versions of resources.
- specifying alternative download mechanisms, such as alternative mirrors.
- send notifications about resource updates.

ResourceSync specifies how to 'enhance' a ResourceSync enabled data source with a notifications mechanism based on WebSub [6]. WebSub specifies the communication between publishers of any kind of Web content and their subscribers, based on HTTP. We further describe WebSub in Section 4.8.

This detailed synchronization information provided by ResourceSync allows for more efficient ways of keeping resources synchronized between a source and a destination than Sitemaps and any other technology that we have analysed.

The extra functionality of ResourceSync over Sitemaps, also increases the technical barriers for its adoption. At the time of writing of this document, we have not yet been able to locate a case of ResourceSync deployment in the cultural heritage domain. Most applications of ResourceSync are in grey literature repositories, which are usually out of scope of cultural heritage.

Since the current focus of Europeana is in acquisition of metadata, ResourceSync may offer more than is necessary, and be an unnecessary challenge for implementation by data providers. Still, ResourceSync is an important technology to follow, particularly as the aggregation of content as well as metadata is starting to gain more attention within the Europeana Network.

4.4. Open publication distribution system

Open Publication Distribution System (OPDS) is a syndication format for digital publications which enables the aggregation, distribution, and discovery of books, journals, and other digital content by any

user, from any source, in any digital format, on any device. The OPDS Catalogs specification [17] is based on the Atom syndication format and prioritizes simplicity. OPDS usage can be found in eBook reading systems, publishers, and distributors. Publishers and libraries have been early adopters of OPDS. We could not yet determine how widely used OPDS is within the Europeana network.

4.5. Linked data platform

Linked Data Platform [15] specifies the use of HTTP and RDF techniques for accessing and manipulating resources exposed as Linked Data [2].

Several cultural heritage institutions publish, as linked data, the metadata regarding their resources. Although linked data publication allows for a standard way to reach the metadata in an automatized way, the standard practices do not address all the requirements needed for metadata aggregation. Mainly two aspects need to be addressed for linked data sources: first, a mechanism for allowing cultural heritage institutions to indicate to aggregators which metadata resources are to be aggregated; and second, an efficient mechanism for allowing efficient incremental harvesting.

Within the many aspects specified by the Linked Data Platform, some provide the necessary standardization for an efficient aggregation based on linked data sources. In particular, the Linked Data Platform Containers and the specification of the usage of HTTP 1.1 could fulfil the requirements for metadata aggregation.

4.6. Webmention

Webmention is a technology that addresses the general problem of allowing Web authors to obtain notifications when other authors link to one of their documents [11]. Webmention is currently published at W3C as a First Public Working Draft. We could not accurately determine how widely adopted Webmention is nowadays, but many resources can be found in the World Wide Web, from software implementations, running services, and many discussions on its use.

The notification mechanisms provided by Webmention, can be used to mediate the communication between the systems of aggregators and the data providers. Webmention presents the following positive aspects:

- A very simple technological solution;
- Any of the parties may initiate the exchange of information.

There are, however, some negative points regarding Webmention:

- No deployments of Webmention are known to exist in cultural heritage institutions;
- The notifications do not allow data to be transmitted, so it must be complemented with other technology, such as the example of linked data, which is described further ahead in this section;
- The notifications may lack semantic meaning (e.g. type of notifications) required for some aggregation operations;
- The application of Webmention, for metadata aggregation, diverges somewhat from what Webmention was designed for. If Europeana uses it for this purpose, further elaboration of specifications will be necessary to define how Webmention is meant to be used.

Due to the lack of a mechanism to transmit data in Webmention notifications, we see its application only in combination with other technologies. For example, in combination with existing linked open data

(LOD) that data providers already have in place. Webmention would allow data providers to indicate to aggregators, which resources from their LOD dataset should be aggregated.

Webmention could also be applied in a similar way to aggregate metadata from IIIF endpoints. The underlying approach may be the same as for LOD. But in this case, the notifications sent by the data providers to aggregators, would contain links to IIIF resources (manifests), and aggregators would use a IIIF crawler to harvest the metadata from the IIIF endpoint.

4.7. Linked data notifications

Linked Data Notifications [3] (LDN) is similar in functionality to Webmention, but it is built having the Web of Data in mind, while Webmention is focused in the Web of Documents. LDN is being designed on top of the W3C's Linked Data Platform (see below), and its notifications have richer semantics than the simple notifications of Webmention. Another promising aspect of LDN is that the notifications may carry data, thus allowing for a more straightforward way of fulfilling metadata aggregation than Webmention. We engaged with the LDN editorial group, and are currently providing feedback to the LDN specifications, considering the metadata aggregation use case.

4.8. WebSub

WebSub [6] provides another option for notification based mechanisms. WebSub specifies the communication between publishers of any kind of Web content and their subscribers, based on HTTP. Subscription requests are relayed through hubs, which validate and verify the request. Hubs then distribute new and updated content to subscribers when it becomes available.

The relevance of WebSub for metadata aggregation comes from its use in ResourceSync (described in Section 4.3). ResourceSync specifies how to 'enhance' a ResourceSync enabled data source with a notifications mechanism.

Although WebSub could be applied as a complementary to other metadata aggregation mechanisms than ResourceSync, such as linked open data or IIIF, we believe that WebSub is the least viable of the three notification based mechanisms we have reviewed. It is mostly unknown in cultural heritage, and in comparison, Webmention or LDN appear to be more advantageous. WebSub does not have a strong industry support and applications in the social web, as Webmention. Compared to LDN, WebSub lacks the semantic aspects of LDN notifications.

4.9. Schema.org

Schema.org [7] was the single new technology we have identified, which could fulfill the requirements of metadata aggregation in the area of data modeling and representation, particularly due to the requirements identified for this area in the Europeana Network.

Schema.org is a cross-domain initiative for structured data in the Internet. Its main application is in web pages, where data can be referenced or embedded in many different encodings, including RDFa, Microdata and JSON-LD. It is developed as a vocabulary (), following the Semantic Web principles. It includes entities and relationships between entities.

Web pages containing Schema.org can be processed by search engines and applications using this structured data, in addition to text and links. The Schema.org website reports its usage in more than 10 million sites and Google, Microsoft, Pinterest, Yandex, among others, already provide services and applications that are based on the available Schema.org structured data. They can, for example, know that a

web page describes a culinary recipe, its ingredients and preparation method, or that it describes a movie, its actors, user reviews, etc. For cultural heritage digital libraries, Schema.org allows the description of books, maps, visual art, music recordings, and many other kinds of cultural resources.

Schema.org is a collaborative and community based activity and its main platform of collaboration is the W3C Schema.org Community Group. The Community Group also serves as a hub for discussion with other related communities, at W3C and elsewhere. Other W3C Community Groups exist that are focused on specific domains, such as health, sports, bibliography, etc. Representatives of the cultural heritage community may be involved this way, should a need to 'improve' Schema.org for cultural heritage aggregation be raised.

It is possible to represent cultural heritage objects using the Schema.org vocabulary. The most relevant classes are schema:CreativeWork and several of its refining subclasses, such as schema:VisualArtwork, schema:Book, schema:Painting, and schema:Sculpture.

Each of these subclasses may be used with more specific properties than the ones available for schema:CreativeWork such as schema:artMedium for schema:VisualArtwork.

The representation of the digital version of cultural heritage objects can also be achieved with schema:MediaObject and its subclasses schema:ImageObject, schema:VideoObject, schema: AudioObject.

Schema.org can also be extended to cover particular cases requiring properties or terms currently not available in the model. These extensions are either approved as part of the core Schema.org or are managed externally. Two of these extensions are of relevance to cultural heritage:

- The Bibliographic Extension provides additional properties and types to describe bibliographic resources. For example, terms such as 'atlas', 'newspaper', 'work and translation', or relationships such as schema:exampleOfWork and schema:workExample.
- The Architypes extension currently works on identifying relevant types and properties to describe archives and their contents. The current proposal defines three new classes: Archive, ArchiveCollection and ArchiveItem.

Designing a metadata aggregation network based in Schema.org may not be a trivial task, due to the large size of the vocabulary. It looks promising in the medium term, however. It has the potential to reduce the effort on data conversion by data providing cultural heritage institutions, since their efforts results in the reuse of the Schema.org data for discovery through Internet search engines and metadata aggregation networks.

5. Conclusion

In conclusion, several technological solutions from the Web are available and look promising for simplifying the implementation of the metadata aggregation scenario in cultural heritage. The next steps of this work will aim to assess the actual usage and existing knowledge of these technologies, within the cultural heritage institutions. Future work, on the technical software side, will address how these technologies may be used for designing crawling robots that aggregate the metadata. We expect that with crawling algorithms, which make use of Web technologies, the technical barriers and operational costs may be lowered, leading to more sustainable metadata aggregation networks.

Acknowledgements

References

[1] Europeana v1.0, The EDM Definition V5.2.7. Online from: http://pro.europeana.eu/web/guest/edm-documentation.

[2] T. Berners-Lee, Linked Data Design Issues. W3C-Internal Document, 2006, available from: http://www.w3.org/DesignIssues/LinkedData.html.

[3] S. Capadisli and A. Guy (eds.), "Linked Data Notifications", W3C Working Draft, 2016, available from: https://www.w3.org/TR/ldn/.

[4] V. Charles and A. Isaac, Enhancing the Europeana Data Model (EDM), Project Europeana V3.0, 2015, available from: http://pro.europeana.eu/files/Europeana_Professional/Publications/EDM_WhitePaper_17062015.pdf.

[5] Digital Public Library of America. Metadata Application Profile, version 4.0, 2015, available from: https://dp.la/info/wp-content/uploads/2015/03/MAPv4.pdf.

[6] J. Genestoux and A. Parecki (eds), WebSub. W3C Candidate Recommendation, 2017, available from: https://www.w3.org/TR/websub/.

[7] Google, Inc., Yahoo, Inc., Microsoft Corporation and Yandex. About Schema.org, available at: http://schema.org/docs/about.html.

[8] S. Gradmann, Knowledge = Information in Context: on the Importance of Semantic Contextualisation in Europeana, 2015, available from: http://pro.europeana.eu/files/Europeana_Professional/Publications/Europeana%20White%20Paper%201.pdf.

[9] C. Lagoze, H. van de Sompel, M.L. Nelson and S. Warner, The Open Archives Initiative Protocol for Metadata Harvesting, Version 2.0, 2002, available from: http://www.openarchives.org/OAI/2.0/openarchivesprotocol.htm.

[10] National Information Standards Organization. ResourceSync Framework Specification, 2014, available from: http://www.niso.org/apps/group_public/download.php/12904/z39-99-2014_resourcesync.pdf.

[11] A. Parecki (ed.), Webmention. W3C Candidate Recommendation, 2016, available from: https://www.w3.org/TR/webmention/.

[12] G. Pedrosa, P. Georg, C. Concordia and N. Aloia, Europeana OAI-PMH Infrastructure, Project Europeana Connect deliverable D5.3.1, 2010.

[13] L. Richardson and S. Ruby, Restful Web Services, O'Reilly, 2007.

[14] Sitemaps XML format. Available from: https://www.sitemaps.org/protocol.html.

[15] S. Speicher, J. Arwe and A. Malhotra, Linked Data Data Platform 1.0. W3C Recommendation, 2015, available from: https://www.w3.org/TR/ldp/.

[16] S. Stuart, R. Sanderson and T. Cramer, The International Image Interoperability Framework (IIIF): A community & technology approach for web-based images, *Archiving* 2015, 2015, available from: http://purl.stanford.edu/df650pk4327.

[17] The openpub community. OPDS Catalog 1.1 specification, 2011, available from: http://opds-spec.org/specs/opds-catalog-1-1.

[18] H. van de Sompel and M.L. Nelson, Reminiscing about 15 years of interoperability efforts, *D-Lib Magazine* **21**(11/12) (2015). doi:10.1045/november2015-vandesompel.

[19] T. van Veen and B. Oldroyd, Search and retrieval in the European library: A new approach, *D-Lib Magazine* **10**(2) (2004). ISSN 1082-9873.