Linked Open Data for Numismatic Library, Archive and Museum Integration

Ethan GRUBER

American Numismatic Society
gruber@numismatics.org

Abstract

The American Numismatic Society (ANS), founded in 1858, is a research institute focusing coins from all eras and regions. It owns one of the largest collections of coins in the world, one of the largest numismatic libraries, is a publisher of monographs and journals, and maintains an archive of research notes from scholars associated with the Society. The Society's collection has been online since 1999, and it has been involved in various typological or hoard cataloguing projects that conform to Linked Open Data methodologies since 2011. These projects have been published at previous CAA conferences and in other venues, but the ANS is now undertaking the digitisation of its archival and library holdings. This paper will focus mainly on the integration of these materials into the ANS's existing data sets to form a cohesive research platform for numismatics that follows open standards and practices established by the broader Library, Archive, and Museum communities.

Keywords: semantic web, linked data, numismatics

ANS Digital Projects to Date

The Collection

The American Numismatic Society's collection comprises more than 700,000 monetary objects (coins, tokens, paper money, etc.) and organisationally similar non-monetary objects that are traditionally associated with the field of numismatics (medals, military decorations). These objects hail from all periods and cultures, but the collection is particularly strong in the Greek and Roman departments. As a result, most of our digital projects are related to these cultures. Digitisation of the collection began in the 1980s with the introduction of a DOS-based curatorial database. The database has gone through several migrations into its current form of FileMaker Pro. The web-based front end, available at http://numismatics.org/search, was first launched in 2011. Called MANTIS, this user interface is built on open source server applications and a middleware framework called Numishare (https://github.com/ewg118/numishare), which is the core application for most of the ANS's digital numismatic projects. FileMaker Pro data are exported as CSV and processed by a programming script into an XML schema, NUDS, which is the underlying data model Numishare is designed to publish. Today, nearly 600,000 objects are available in MANTIS, and more than 20% have been photographed.

Nomisma.org

Contemporaneous to the initial release of MANTIS was the development of Nomisma.org, an open access thesaurus of numismatic concepts that conforms to the principles of Linked Open Data (LOD). Many of the attributes by which coins are traditionally classified—denomination, material, and production places ('mints'), etc.—are reusable across many coins, which necessitates the publication of unique identifiers to define them. These identifiers are URIs, following LOD methodologies, where human- and machine-readable information may be accessed. The mint "Rome" is defined by http://nomisma.org/id/rome, and the Roman emperor Augustus, is defined by http://nomisma.org/id/augustus. A user may see a variety of information (maps, labels, matching URIs, example coins) by visiting one of these URIs in a browser, but the user may also request RDF conforming to a variety of serialisations via REST or HTTP content negotiation.

Initially a proof of concept created in 2010 by Andrew Meadows and Sebastian Heath, then both of the American Numismatic Society, Nomisma has grown into a widely recognised web service, guided by an international scientific committee of numismatists and information/computer scientists. Its usage in the cultural heritage sector has increased steadily over the years (Smith-Yoshimura 2016). There is much to be said about Nomisma's evolution into the backbone for the future of numismatics, but these topics have been covered in various venues, both in terms of high-level vision (Gruber *et al.* 2013b, Meadows and Gruber, 2014) and low-level information architecture (Gruber, 2016) or mapping relational databases into RDF ontologies (Tolle and Wigg-Wolf, 2016).

It was apparent from the development of MANTIS that integration of Nomisma.org data would be vital to drive controlled vocabulary for faceted browsing and

geospatial mapping of mints. The NUDS XML schema was adapted to allow the linking of objects to URIs defining intellectual concepts published on Nomisma. org. Records in MANTIS may also link to URIs that define published typologies or coin hoards (often defined as three or more coins found in the same archaeological context). These databases are discussed below.

Coin Typology Databases

Coins have historically been categorised by a variety of individual attributes: the manufacture process, material, monetary denomination, mint, date, entities responsible for issuing the coin (whether individual rulers or corporate organisations), and the iconography and inscriptions (or 'legend' in numismatic terminology) on the front and back (obverse and reverse) of the coin. The combination of each of these individual attributes comprises a coin 'type,' and types are often uniquely numbered, thematically organised, and published in volumes of printed books.

For example, Roman Imperial coins have been published in numerous volumes over the last several centuries, but the standard reference work today is the ten volume Roman Imperial Coinage (RIC). Collections of Imperial coins therefore refer to standard type numbers from RIC, e.g., Augustus 1a, a silver denarius minted in Emerita (Spain) from 25–23 B.C. These numbers were once printed in collection inventories or cards associated with each coin, but are now inserted into bibliographic fields in museum databases.

In 2012, the ANS published the first edition of Online Coins of the Roman Empire (OCRE: http://numismatics. org/ocre/), a digital type corpus based on the RIC numbering system. Following the patterns established by Nomisma.org, URIs were created for each RIC type number. RIC Augustus 1a is represented by http:// numismatics.org/ocre/id/ric.1(2).aug.1A, where the 'ric.1(2)' in the ID sequence refers to RIC volume 1, second edition (Sutherland and Carson, 1984). OCRE was discussed in Gruber *et al.* (2013a), and while the functionality remains similar as it did in 2012, the architecture has evolved considerably to integrate more Semantic Web technologies, such as SPARQL.

The introduction of the Nomisma.org SPARQL endpoint in early 2013 facilitated a broader incorporation of materials from Nomisma partner institutions. The Berlin Münzkabinett was the first partner to make its Roman Imperial coins available in OCRE, following by large collections like the British Museum and smaller ones like The Fralin Museum at the University of Virginia. Coins have recently been incorporated from finds and archaeological databases, such as the UK's

Portable Antiquities Scheme and OpenContext.org, managed by the University of California-Berkeley.

In 2015, the scope of online typologies was extended from Roman Imperial to Republican coinage with the release of Coinage of the Roman Republic Online (CRRO: http://numismatics.org/crro/) and PELLA (http:// numismatics.org/pella/), the coinage issued in the name of Alexander the Great. The ANS is working to expand PELLA to include all Macedonian coinage, and will be publishing type corpora for the other Hellenistic kingdoms in the coming years. In total, about 15 partner institutions are contributing data for about 100,000 physical specimens associated with the type URIs published in OCRE, CRRO, and PELLA.

Coin Hoard Databases

The ANS's first online coin hoard project launched in 2013 as a collaboration with Kris Lockyear at University College London. This database, transformed from Lockyear's personal Microsoft Access research database into NUDS, was published in Numishare as Coin Hoards of the Roman Republic (CHRR: http://numismatics.org/ chrr/) (Gruber and Lockyear, 2015). Building on previous linked data work, CHRR extracts typological data from CRRO in real-time for display and analysis purposes. Each hoard is available at a URI, and these URIs have been incorporated into databases of Nomisma partner institutions, such as Berlin, which has made it possible to map find spots within their own collection.

In the Greek realm, the American Numismatic Society has published multiple print volumes of information and interpretation of coin hoards. An Inventory of Greek Coin Hoards (IGCH) (Thompson *et al.*, 1973) includes information for nearly 2,400 Greek coin hoards, culled from publications, personal research notes, and other archival materials. Following IGCH, ten supplemental Coin Hoards volumes have been published. A prototype of the IGCH data is available at http://coinhoards.org, but much work remains to be done to make this project as robust as Coin Hoards of the Roman Republic. Nevertheless, the URIs published by this system are stable and have been incorporated into both MANTIS and the Berlin Münzkabinett, which has already opened the door to the direct linking between physical specimens, the typologies they represent, and the patterns of their geographic circulation.

The integration of these materials can be seen clearly at http://nomisma.org/id/tetradrachm, the URI that defines a Greek denomination. The map displays layers for the mints that produced tetradrachms (essentially a web service that transformations a SPARQL query into geoJSON, derived from types published in PELLA) and the places where they have been found (mainly derived
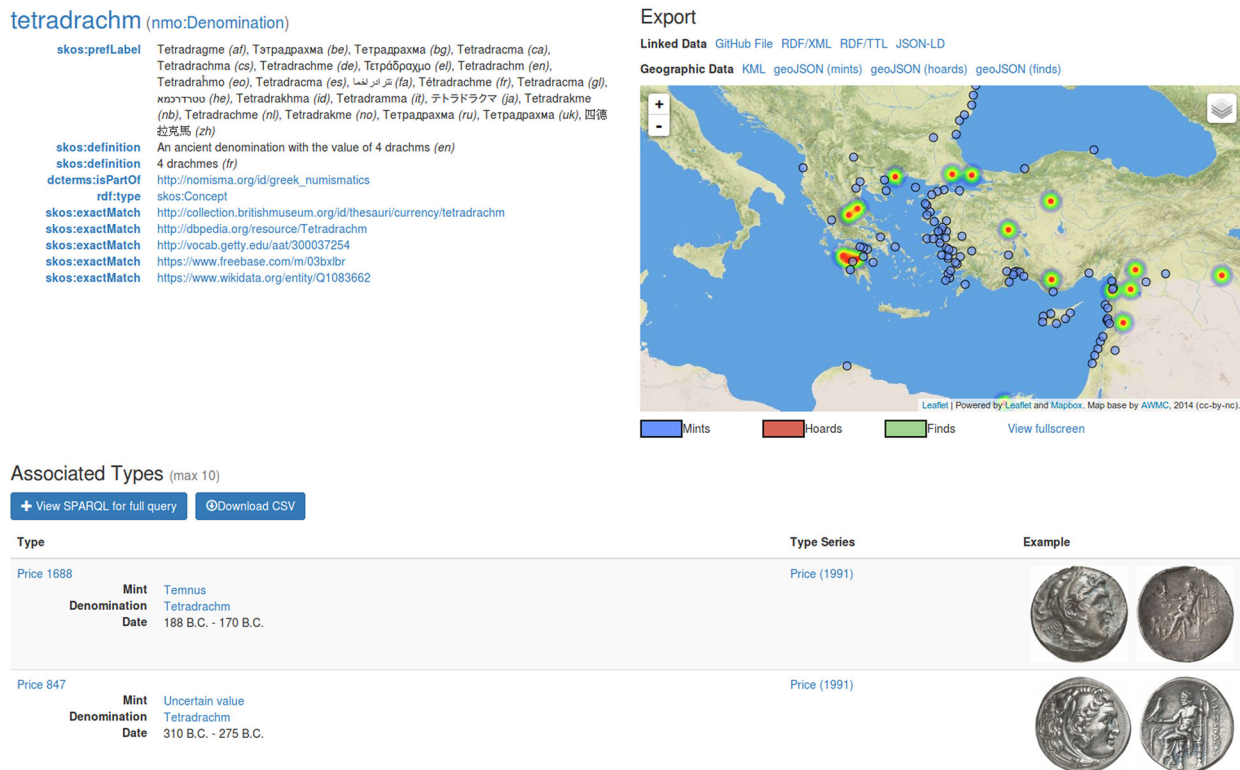
Figure 1. Screenshot of http://nomisma.org/id/tetradrachm

from find spots in IGCH). Below the map are examples of related typologies, including photographs of coins of these types (Figure 1). The user may click on a hoard in the map to go to IGCH 1670 (http://coinhoards.org/id/igch1670), and similarly, there are examples of coin types that have been found within this hoard, derived from MANTIS and Berlin. Additionally, there may be links to research notebooks or digitised monographs, as will be discussed below.

The Archives

Shortly following the release of the first version of MANTIS in spring 2011 came the release of the American Numismatic Society digital archive, Archer (http://numismatics.org/archives). In its first stage, Archer was an editing and publication framework for finding aids, which are documents that provide a description of the contents of an archival collection. This description may include the hierarchical organisation of the materials into boxes and folders, down to the level of an individual item. Finding aids were traditionally paper, but are now typically encoded in an XML standard called Encoded Archival Description (EAD). Archer is built on EADitor (https://github.com/ewg118/eaditor), an open source middleware framework for the creation and publication of EAD finding aids. Architecturally similar to Numishare, EADitor was expanded to support the publication of other metadata standards common to

the Library and Archives sectors, such as the Metadata Object Description Schema (MODS), a bibliographic model, and the Text Encoding Initiative (TEI), which is a mark-up language for textual transcription.

Like Numishare, EADitor has evolved considerably to be more LOD-aware, both in terms of web service lookup mechanisms (to link people to authority files published by the Virtual International Authority File (VIAF.org), link places to Geonames.org or the Pleiades Gazetteer of Ancient Places, and link genres/formats to the Getty Art and Architecture Thesaurus) and in the serialisation of EAD, MODS, and TEI into RDF to facilitate the interlinking between archival resources.

With a grant of $7,500 from the Gladys Krieble Delmas Foundation, in 2014 the ANS digitised 43 notebooks by prominent Greek numismatist and former ANS President, Edward T. Newell. This project entailed the generation of a Text Encoding Initiative (TEI) XML file for each notebook, containing bibliographic metadata extracted from the ANS's library catalogue DONUM and a list of links to the scanned page images. By extending the functionality of EADitor to support TEI publication and image annotation with Annotorious (http://annotorious.github.io/), Librarian David Hill and Assistant Archivist Arnold Tescher were able to rigorously annotate the coins, hoards, scholars, bibliographic publications, geographic places, and
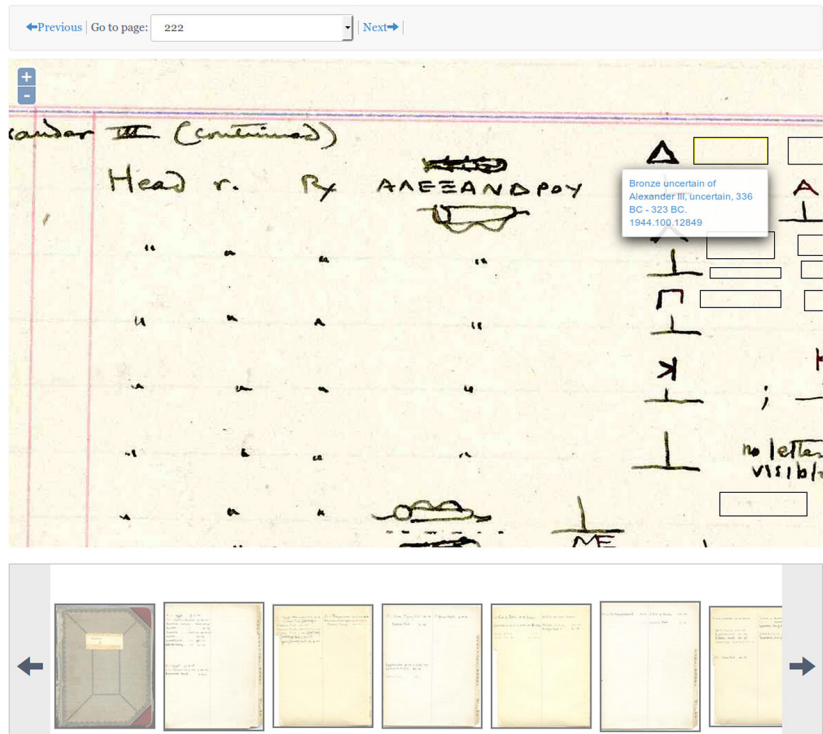
Figure 2 - The user interface of an annotated notebook in EADitor

other sorts of LOD-defined entities in a handful of these notebooks (Figure 2). There remains much more work to be done to complete the annotation project.

A fundamental aspect of the ANS Archives is its personal and corporate authority system, http://numismatics.org/authorities/. Publicly launched in 2014, the ANS Authorities section of Archer includes records for more than 100 people and corporate bodies associated with the Society. These include individuals that were members or officers of the Society or were prominent scholars in the field, for whom the ANS maintains archival materials (e.g., correspondences). These authority files are authored in another standard XML schema from the archival community, Encoded Archival Context—Persons, Corporate Bodies, Families (EAC-CPF), and published in the open source xEAC application (https://github.com/ewg118/xEAC), which began development in 2012.

EADitor and xEAC both function as standalone applications, but can be configured to interoperate. That is to say, EADitor may interact with the Atom feed API published by xEAC in its EAD editing back-end to incorporate URIs minted by xEAC directly into finding aids. Both applications may be configured to interact with a SPARQL endpoint to create, update,

and delete triples associated with archival authorities and materials. Therefore, a finding aid detailing the collection of Edward T. Newell (http://numismatics.org/archives/ark:/53695/nnan0084) or a research notebook (http://numismatics.org/archives/ark:/53695/nnan187715) will include the URI for Newell in the ANS Authorities system (http://numismatics.org/authority/newell) as the creator (dcterms:creator). When publishing the EAC-CPF file in xEAC, the XML file will be transformed into RDF following a variety of ontologies relevant to biographical information and posted into the triplestore via the SPARQL/Update protocol. Likewise, EAD finding aids and TEI notebooks will be serialised into RDF conforming to Dublin Core Terms, Open Annotation, and other common ontologies and published from EADitor into the SPARQL endpoint. When one visits the URI for Edward Newell, that person may get a list of archival and library materials held by the ANS or other institutions. As a result, the page for Newell is more than a collection of biographical information and related people and organisations, but a gateway to more information by or about the scholar. In a few short clicks, a researcher might traverse the graph of numismatic information from Newell's biography to a notebook in the ANS archives to an IGCH hoard to a coin type in PELLA to a coin in MANTIS to the Nomisma.org URI for tetradrachm, which may include

other, related research materials hosted by external services. The research potential is enormous, as scholars do not need to invest as much time acquiring information as before, and may devote more resources into the interpretation of evidence.

The Library

The American Numismatic Society maintains the largest collection of specialist numismatic literature in the world. Many patrons physically travel to the ANS office in New York City every year to conduct research, as many materials (like old auction catalogues or esoteric numismatic series) are very rare and generally inaccessible to the public. Among the materials in the ANS library are series and monographs published by the Society itself. As it has demonstrated in other projects detailed in this paper, the ANS is fully committed to the free and open proliferation of knowledge, and it is striving to make its own out of print (but often, still in copyright) content openly accessible to the public.

As part of the Google Books project, roughly 500 monographs were scanned by Google's academic library partners. The scans of these monographs were transferred to HathiTrust for long-term sustainability, but only those published before 1923 are in the public domain. We have since issued a Creative Commons license for all works hosted by HathiTrust, making every one of them open. Building on this increased accessibility, the ANS applied for and received a grant for approximately $50,000 from the Mellon Foundation in December 2015 to transcribe about 80 of the rarest works into TEI, enhance with mark-up linking coins, hoards, entities, etc. mentioned in the books to URIs in our or other databases, and make these works available as EPUB 3.0.1 documents.

All of these monographs have been transcribed into TEI, but only a few are available online through the ANS Digital Library (http://numismatics.org/digitallibrary/), as of the date of this publication. The remaining will be completed by the end of 2016. The Digital Library is built on ETDPub (https://github.com/AmericanNumismaticSociety/etdpub), which was initially conceived as a framework for publishing numismatic electronic theses and dissertations (including the full-text indexing of PDFs and LibreOffice/Word document files into Apache Solr). In early 2016, ETDPub was extended to support the publication of TEI and serialisation of TEI into EPUB, which is built upon the XML Pipeline Language inherent to Orbeon XForms' architecture. Like EADitor, monographs may be associated with URIs published in the ANS Biographies, and TEI files are transformed into RDF conforming to Open Annotation. As a result, researchers may access Coin Hoards (http://numismatics.org/digitallibrary/
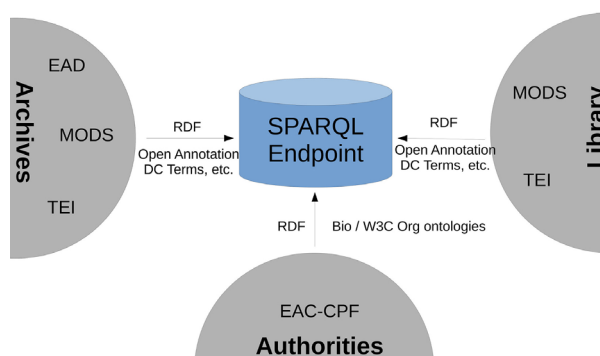

Figure 3 - Interoperable models

ark:/53695/nnan146115) from the biography of Sidney Noe (http://numismatics.org/authority/noe), which will provide access to the Saida Find (IGCH 1508) (http://coinhoards.org/id/igch1508). This hoard includes a reference to a specific page in one of Newell's research notebooks as well as points back to the section in Noe's book about the hoard.

What once began simply as an online database of the ANS's numismatic collection has evolved into a suite of purpose-built tools for publishing coins, types, hoards, archives, and library holdings, which are becoming increasingly interlinked—not just internally with respect to the Society's materials—but also paving the way these materials to be made available to scholars through other external services.

The Underlying Architecture

Application Stack

From the above section, one may surmise that ANS digital projects follow a similar architectural pattern, which is mainly a product of the author's software development experience gained from the academic library realm, which is highly dependent upon XML technology. Numishare and EADitor were both borne from the Scholars' Lab at the University of Virginia between ca 2008–2010. These frameworks have advanced in a myriad of ways over the last several years, but three open source Java-based server applications remain fundamental aspects of the architecture of each framework. ETDPub and xEAC are built on this stack, and are designed to handle the creation, management, and publication of other types of documents. The server applications, which run in Apache Tomcat, are as follows:

Apache Solr (http://lucene.apache.org/solr/): Enables full text searching and faceted browse

eXist XML database (http://exist-db.org): Since many of the models that form the basis of the ANS digital

collection, archives, and library are XML schemas, we use the eXist XML database. The eXist database supports XQuery, a query language for XML, which is used for back-end search and publication functionality, as well as for some APIs in Numishare.

Orbeon XForms (http://www.orbeon.com): The public user interfaces of these projects are driven by the Orbeon XML Pipeline Language (XPL), a data pipeline system for implementing Model-View-Controller (MVC) architecture in an XML framework. These pipelines enable the transformation of various XML or JSON data models (whether Solr search results, XML documents in eXist, SPARQL responses, or external REST services) into a wide variety of serialisations, from HTML for use in browsers to geoJSON or KML for geographic visualizations to linked data in the form of RDF/XML, Turtle, and JSON-LD. The back-end of these applications is driven by XForms, a W3C specification for advanced web form and data processing functionality.

Apache Fuseki (https://jena.apache.org/documentation/serving_data/): A triplestore and SPARQL endpoint produced by the Jena project. It was chosen over other endpoints for its ease of use and excellent documentation. First tested in December 2012, it was placed into production in spring 2013 when OCRE was reengineered to interact with data aggregated with Nomisma.org's SPARQL endpoint. A second Fuseki triplestore is hosted by numismatics.org for aggregating library and archival linked data. Any SPARQL 1.1 compliant endpoint may be deployed in this application stack.

The development of these projects has been gradual over the last five years, with new functionality and enhanced interoperability introduced in iterative phases. MANTIS, which was once a data silo, became more incorporated with the broader Linked Open Data cloud through the integration of URIs published by Nomisma.org. With the launch of coin typology and coin hoard projects from 2013–2015, each project became more interactive, with data constantly flowing from one project to another. Nomisma itself is the keystone for the future of the discipline. The aggregation of physical specimens related to typologies, the metrical analyses of weights and diameters as part of the OCRE, CRRO, and PELLA interfaces, and the wide variety of geographic visualizations are all driven by the SPARQL endpoint hosted by Nomisma.org. SPARQL queries are so intertwined into the many user interfaces of these projects that the advanced Semantic Web information system underneath is invisible.

Much like MANTIS, Archer once was a silo of archival material—primarily EAD finding aids—that became increasingly interlinked with external authority systems like VIAF, the Library of Congress Subject Headings, and the Getty vocabularies. With the introduction of the ANS Authorities project, Semantic Web technology opened the door to linking between our own archival materials with internally-managed personal and corporate authorities. Even still, this integrated archives/authority system formed a kind of silo in its own right. The deployment of the ANS Digital Library extended the scope of our triplestore, enabling the linking between archival materials, digitised books, and our archival authority system.

The information about coin types, hoards, and physical specimens hosted by museum or archaeological database systems forms a powerful numismatic data ecosystem, but these databases form only a portion of the total human knowledge base regarding the study of coins. The ANS Library and Archives are another portion of this knowledge base, and much effort has been placed at the American Numismatic Society over the last year to better integrate the numismatic data ecosystem with the numismatic document ecosystem into a cohesive research platform. The publication of Edward Newell's research notebooks and electronic monographs as part of the NEH/Mellon Humanities Open Book Program have been instrumental in developing and testing a more thorough integration of these data and documents, as these documents include references to coins in the ANS and other collections as well as types and hoards published in digital catalogues like OCRE and IGCH.

Semantic Web Technology for Complete LAM Integration

The data models and software architecture underlying the ANS and Nomisma.org numismatic data ecosystem have been well published, even as recently as the 2016 CAA conference in Oslo (Tolle and Wigg-Wolf, forthcoming) and XML London (Gruber, 2016), but it is nevertheless useful to briefly outline how this system presently functions.

As mentioned above, the ANS digital projects published in Numishare are encoded in NUDS/XML. An RDF export feature is inherent to Numishare's functionality; essentially it is a workflow that executes XQuery to aggregate NUDS in the eXist database to pipe through XSLT into RDF. This RDF conforms to the Nomisma ontology (http://nomisma.org/ontology), but also incorporates other common ontologies such as Dublin Core Terms, Friend of a Friend (FOAF), and others. The precise model varies based upon the class of data object. Documentation for contributing data to Nomisma is available at http://nomisma.org/documentation/contribute.

RDF for physical specimens from MANTIS and partner institutions like the Portable Antiquities Scheme and Berlin Münzkabinett, as well as hoard projects like CHRR, link to type URIs published by OCRE, CRRO, and PELLA. The coin type RDF links to URIs published by Nomisma. All of these data sets are aggregated into a central SPARQL endpoint, enabling semantic reasoning across coins and hoards. A physical coin does not need to be explicitly designated as a 'denarius' when it links to a coin type that contains the property nmo:hasDenomination (in the Nomisma ontology) that links to http://nomisma.org/id/denarius.

Nomisma itself hosts a number of APIs that are queried by Numishare, such as 'Average Weight', which delivers a numerical response to a short-hand SPARQL query executed with request parameters (e.g., the average weight of denarii of Augustus from AD 10–15). Other more complex queries (for geographic visualization, for example) are executed directly in Numishare's XPL models and transformed through XSLT into KML or JSON for display by open source Javascript mapping libraries, like OpenLayers or Leaflet.

The document ecosystem of the American Numismatic Society is fairly similar to that of the numismatic data system, since the implementation of the EADitor–xEAC interactivity functionality was influenced by the lessons learned in the development of Nomisma.org and the aggregation of numismatic content. The difference with the implementation of Linked Open Data methodologies to archival materials is that there are no true standards for representing archival collections as a graph of information. Experiments in mapping archival collections to CIDOC-CRM have been conducted (Hennicke, 2013, Halling, 2016), but particular models of the CRM may vary from institution to institution. There is no one way to model anything in CIDOC-CRM, nor, arguably, should there be. EAD, like TEI, is a document model, and it is simply impossible to map all content from a document model into a graph. This is not the intention of linked data. However, digital books and archival collections may be represented as a graph in those areas that pertain specifically to linking web resources together: a) metadata (authors, genres, subject headings, dates of publication, etc.), and b) aspects of the body of the document that may be linked to other information systems.

For lack of a purpose-built archival Conceptual Reference Model, we chose to map document metadata (for electronic monographs, archival notebooks, finding aids, and other archival materials) to existing, commonly used ontologies on the web, such as Dublin Core Terms. We implemented Open Annotation for linking resources mentioned in sections of the body of a document. Furthermore, we used Dublin Core Terms to hierarchically link the components of documents together: sections and chapters in TEI files and series, sub-series, folders, and items in EAD finding aids. This would allow users to navigate from an object in MANTIS to the lowest level section of an EBook published in our Digital Library (see http://numismatics.org/collection/0000.999.20456, for example) or directly to the page of one of Newell's annotated notebooks (http://numismatics.org/collection/1944.100.12599). The RDF from our library, archives, and authority system are aggregated into a single SPARQL endpoint, which can be queried by the ANS Authorities project, MANTIS, IGCH, or any external system. After years of development, we are now able to begin connecting our numismatic data and document systems together to form a comprehensive research platform, a suite of tools that will continue to grow in its usefulness as more data and documents are incorporated into the system. Linked Open Data principles are the foundation for this system.

Conclusion and Future Work

We have clearly demonstrated the potential for Linked Open Data methodologies to enhance access to numismatic materials as well as bind these materials more seamlessly together into a sophisticated research framework that incorporates museum and archaeological collections, hoard information, typologies, archival research materials, and digitised auction catalogues, journals, and monographs. Tens of thousands of Greco-Roman typologies have been published online, along with thousands of hoards, and data for 100,000 physical specimens have been aggregated, interlinked with these hoards and typologies. Within the next year, nearly 100 rare and out-of-print monographs will be openly accessible, annotated with references to coins, types, hoards, and a wide variety of other resources. So while there are currently a relatively small handful of examples of objects in the ANS numismatic collection link to books and archival materials, there will be a significant increase in interlinked content within the next few years as the Humanities Open Book Program and the Newell annotation projects draw to a close.

Nomisma.org is central to the success of these projects. One of Nomisma's greatest advantages is the effort placed by the scientific committee in co-referencing between Nomisma URIs in concepts in other authority systems. One of these is the Pleiades Gazetteer of Ancient Places. By annotating mints and regions mentioned in our textual materials with Nomisma URIs, we are able to expose content to Pelagios Commons. Already, several ebooks are accessible in Peripleo (http://pelagios.org/peripleo/map), making them available to researchers interested in ancient geography. We anticipate that

other external aggregation projects will arise, and we hope that we might be able to make our material more broadly accessible to researchers who might otherwise not be aware of our content. Serendipity in research is one of the great potentials in developing large-scale Linked Open Data systems within the Library, Archive, and Museum communities.

Many different communities within cultural heritage and academics are working diligently to interlink their own disciplines, while simultaneously working to bridge gaps between fields. We are confident that in the coming decades, the numismatic research ecosystem will become more interwoven into the fabric of human knowledge.

References

Gruber, E. (2016) 'Using XForms to create, manage, and publish Linked Open Data', in *XML London 2016 Conference proceedings. University College London, London, United Kingdom, June 4-5, 2016.* XML London, pp. 99–107. doi:10.14337/XMLLondon16.Gruber01.

Gruber, E., Bransbourg, G., Heath, S., and Meadows, A. (2013a) 'Linking roman coins: current work at the American Numismatic Society', in Earl, G., Sly, T., Chrysanthi, A., Murrieta-Flores, P., Papadopoulos, C., Romanowska, I., and Wheatley, D. (eds) *Archaeology in the Digital Era. Papers from the 40th Annual Conference of Computer Applications and Quantitative Methods in Archaeology (CAA), Southampton, 26–29 March 2012.* Amsterdam: Amsterdam University Press, pp. 249–258.

Gruber, E., Heath, S., Meadows, A., Pett, D., Tolle, K., and Wigg-Wolf, D. (2013b) 'Semantic web technologies applied to numismatic collections', in Earl, G., Sly, T., Chrysanthi, A., Murrieta-Flores, P., Papadopoulos, C., Romanowska, I., and Wheatley, D. (eds) *Archaeology in the Digital Era. Volume II. e-Papers from the 40th Conference on Computer Applications and Quantitative Methods in Archaeology Southampton, 26–30 March 2012.* Amsterdam: Amsterdam University Press, pp. 264–274. Available at: http://arno.uva.nl/document/545855 (Accessed: 3 March 2017).

Gruber, E. and Lockyear, K. (2015) 'From dBase III+ to the semantic web: twenty-five years of the coin hoards of the Roman Republic database', in Traviglia, A. (ed.) *Across Space and Time. Papers from the 41st Conference on Computer Applications and Quantitative Methods in Archaeology. Perth, 25–28 March 2013.* Amsterdam: Amsterdam University Press, pp. 336–346.

Halling, S. (2016) 'Archival information + CIDOC CRM = true?', *Digisam.* Available at: http://www.digisam.se/index.php/hem/entry/archival-information-cidoc-crm-true (Accessed: 21 July 2016).

Hennicke, S. (2013) 'Representation of archival user needs using CIDOC CRM', in Alexiev, V., Ivanov, V., and Grinberg, M. (eds) *Proceedings of the Workshop Practical Experiences with CIDOC CRM and its Extensions (CRMEX 2013), Valetta, Malta, September 26, 2013.* pp. 48–60. Available at: http://ceur-ws.org/Vol-1117/ (Accessed: 3 March 2017).

Meadows, A. and Gruber, E. (2014) 'Coinage and numismatic methods. A case study of linking a discipline', in Elliott, T., Heath, S., and Muccigrosso, J. (eds) *Current Practice in Linked Open Data for the Ancient World.* ISAW, New York University. Series: ISAW Papers 7.15. Available at: http://dlib.nyu.edu/awdl/isaw/isaw-papers/7/meadows-gruber/ (Accessed: 21 July 2016).

Smith-Yoshimura, K. (2016) 'Analysis of international Linked Data survey for implementers', *D-Lib Magazine*, 22 (7/8). doi: 10.1045/july2016-smith-yoshimura.

Sutherland, C.H.V. and Carson, R.A.G. (eds) (1984) *The Roman Imperial Coinage* (vol. 1). 2nd edn. London: Spink and Son Ltd.

Thompson, M., Mørkholm, O., and Kraay, C.M. (eds) (1973) *An inventory of Greek coin hoards* (vol. 1). New York: American Numismatic Society.

Tolle, K. and Wigg-Wolf, D. (2016) 'How to move from relational to 5 Star Linked Open Data —a numismatic example', in Stefano, C., Scopigno, R., Carpentiero, G., and Cirillo, M. (eds) *CAA2015. Keep the Revolution Going >>> Proceedings of the 43rd Annual Conference on Computer Applications and Quantitative Methods in Archaeology* (vol. 1). Oxford: Archaeopress, pp. 275–281.

Tolle, K. and Wigg-Wolf, D. (forthcoming) 'An ontology for a numismatic island with bridges to others', in Uleberg, E. and Matsumoto, M. (eds) *CAA2016. Exploring Oceans of Data. Proceedings of the 44th Conference on Computer Applications and Quantitative Methods in Archaeology.* Oxford: Archaeopress.