

Linking Roman Coins: Current Work at the American Numismatic Society

Ethan Gruber

American Numismatic Society

Gilles Bransbourg, Sebastian Heath

New York University

Andrew Meadows

American Numismatic Society

Abstract:

This paper details a current project at the American Numismatic Society: Online Coins of the Roman Empire (OCRE), an open-access catalogue of Roman Imperial coin types. The paper discusses linked data approaches applied to numismatic collections, the application architecture of the project, and the practical functionality of the project's web interface, illustrating its value to researchers.

Keywords:

Linked data, Information Architecture, Numismatics

1. Introduction

This paper describes a series of related initiatives to publish Roman coins—the basic material of the discipline of Roman numismatics that works to place coins in their historical and economic context—on the Internet using practices and methods that draw on the “Linked Data” approach to access and re-use of internet-based resources. The components of the work described here are: Nomisma.org, a digital resource that establishes stable URIs for numismatic concepts, with a current focus on both Greek and Roman coins; Numishare, a software platform for the management of numismatic data that supports linked data approaches; Online Coins of the Roman Empire (OCRE), a project recording and publishing all the known types of coins issued by the Roman Empire, accessible at <http://numismatics.org/ocre/>. While all of these projects are publicly available, they continue to be under active development. In combination they demonstrate the power of stable links between high-quality structured data.

In its capitalised form - that is as a defined set of practices - “Linked Data” is “a term used to describe a recommended best practice for exposing, sharing, and connecting pieces of data, information, and knowledge on the Semantic Web using URIs

and RDF.”² In this paper we focus on three aspects of Linked Data and describe how we implement them within the field of Roman coinage. Those three aspects are: stable URIs for each component of the study of coins, meaning URIs for descriptions of coin types, and for the vocabulary terms we use to describe those coins; availability of machine parsable data via those URIs; and rich re-use and interlinking between separate datasets that allows both cross-project efficient and discovery.

Coinage, as one of the ancient world's few examples of a mass-produced medium, is an excellent disciplinary laboratory—so to speak—in which to implement the concepts of Linked Data. Variations such as obverse and reverse designs and inscriptions, together with materials, denominations and dimensions offer a categorisation system that has long been employed by numismatists. It has achieved its canonical form for the coinage of the Roman Empire in the type-corpus Roman Imperial Coinage. This series of ten volumes, originally published between 1923 and 1994 and in the process of being updated, attempts to provide a unique identification number for every recorded variety of Roman Imperial coin, and has come to serve as a standard cataloguing resource for archaeologists and museum curators.

Corresponding author: ewg4xuva@gmail.com

² “Linked Data – Connect Distributed Data across the Web,” [Linkeddata.org](http://linkeddata.org), accessed October 31, 2012, <http://linkeddata.org/>.

Roman Imperial Coinage (RIC) is at one level an internally Linked Data product in its own right: it employs permutations of stable numismatic identifiers to assert a series of permanent identifiers (type numbers). Furthermore, the complex of information that makes up a coin type, as well as the coin type itself are inherently linkable to other categories of information about the ancient (or indeed modern) world: emperors, cities, findspots, magisterial offices all have significance beyond the world of numismatics. All of these attributes can be assigned stable URIs that facilitate cross-project and machine readable representations of this information through Nomisma.org.

2. Nomisma.org

Nomisma.org - begun in 2010 and hosted by the American Numismatic Society - is an ongoing project to establish stable URIs for numismatic concepts. It is currently available at the URI <http://nomisma.org>. What do we mean by numismatic concepts and why do such concepts need a separate namespace for their definitions? Numismatics has developed since the 16th century as an increasingly formal discipline encompassing the study of coinage and other physical monetary instruments used since the 7th century BC around the world. It is an important feature of numismatics that the majority of coins - here including all metal monetary instruments produced by a political entity to an official weight and with an established and distinctive visual appearance - can be described using a common vocabulary that is applicable to a vast range of ancient and modern examples. While the extremely broad scope of the discipline means that all general statements have exceptions, coins usually have an obverse and reverse, for which the colloquial English terms are “heads” and “tails”. Furthermore, there is only occasional ambiguity about how these two terms are used in the description of a particular coin. More specialised terms include “Legend”, meaning any words - whether abbreviated or not - on a coin, and “Axis” the angular relationship between the obverse and reverse. Most modern coins have an “Axis” of 6 on a scale of 12 because the reverse image is upside down in relation to the obverse. Additionally, it is important to know whether a visual or textual element - the portrait of a ruler or the phrase “E Pluribus Unum” - appears

on the obverse or the reverse. Because these terms have a specialized and well-established meaning, it is useful for the discipline to have a single resource which defines those meanings, both for use within the numismatics, and for the purpose of linking to other fields of study. Accordingly, Nomisma.org is inspired by the desire to promote interoperability, re-use and transparency within and beyond the very specialized field that it addresses.

An example Nomisma.org URI drawn from Roman numismatics is [http://nomisma.org/id/ric.1\(2\).aug.1a](http://nomisma.org/id/ric.1(2).aug.1a). From the perspective of the role of Nomisma within the discipline, the most important feature of that URI is that it is stable, relatively short, and semantically clear within its numismatic context. In particular, the opening characters of the final component of the path, “ric”, are recognizable as the de facto abbreviation for the series Roman Imperial Coinage, the standard reference work establishing identifiers for coins struck under the authority of the Roman Empire. Other segments of this identifier represent volumes (1(2) = volume 1, second edition), emperors (aug = Augustus), and reference numbers (no. 1a). This URI is itself too new to have seen widespread adoption, but it suggests a future in which multiple numismatic collections point that URI as an unambiguous identifier for coins of that type. This is a solution to the current problem whereby existing catalog entries—both online and in print—use various ad hoc abbreviation systems for RIC coins. These can be thought of as labels, and Nomisma does not mandate that collections use any particular sequence of characters for such human-readable indications of RIC type. Pointing to a URI is a sufficient and interoperable indication of shared identity across multiple collections.

It is a principle of Linked Data that there be machine parsable data available via the URIs that identify resources, with a strong preference for an RDF-based serialization. Nomisma.org has adopted XHTML+RDFa 1.1 as the archival representation of the information it represents about each numismatic concept it describes. Figure 1 shows that the RDFa markup is exposed to human readers in the form of labels on the attributes of that coin type. For example, the value “rome” is marked as being the mint. An RDFa distiller, such as that deployed by the W3, will produce the triple: ‘nm:ric.2.tr.432

nm:mint nm:rome' . Note that 'rome' is actually a reference to the resource <http://nomisma.org/id/rome> , this RDFa based resource in turn points to the Pleiades reference for Rome, which is a widely recognized identifier for the ancient capital of the Roman Empire. In this way the Nomisma.org ID facilitates bi-directional links to non-numismatic efforts to establish stable URIs for scholarly concepts.

Given our focus on cross-project reuse of vocabularies and data, this paper moves from describing our efforts to create stable URIs for numismatic concepts, to an application of such concepts within a further specialized field of numismatics, to a description of the software platform that supports specialized applications, and which is currently the basis for hosting and managing further numismatic data. This last point suggests that we are at a point of rapid progress in applying linked data concepts to many areas of the very wide field of numismatics.

3. Numishare: Technical Underpinnings

OCRE has three basic purposes. It is intended first as a typological catalogue of Roman Imperial Coinage that harnesses basic information from the Nomisma.org IDs to make the knowledge encapsulated in the RIC print volumes available to a broad audience, within the archaeological discipline, but also to a broader public. Using the principles of Linked Data it is also intended to provide links from numismatic descriptions to other online resources describing the ancient world, such as, for example the Pleiades project ([\[pleiades.stoa.org/\]\(http://pleiades.stoa.org/\)\). Thirdly, OCRE is designed to be able to accumulate examples of multiple examples of typologically similar coins to facilitate quantitative and qualitative analysis of the material, including metrological, compositional and die-study. To date, the coins from the American Numismatic Society's collection have been added. This was possible due to the prior digitisation of the collection in the Society's MANTIS database \(<http://numismatics.org/search/>\), the curatorial staff of the ANS having manually added all RIC-related information. In due course, other collections will follow. OCRE has been built on the Numishare platform.](http://</p></div><div data-bbox=)

On a technical level, OCRE is built on a data model in which numismatic metadata is represented as XML. Apache Cocoon is the server application for serializing and delivering data. Cocoon is an open-source Java-based framework designed to effectively process XML. In addition to Cocoon, the stack includes other open-source Java-based applications which run in Apache Tomcat on the server: Apache Solr for faceted searching, eXist XML database, and Orbeon XForms for editing, managing, and publishing XML. These applications, including various XSLT stylesheets, Javascript files, and CSS stylesheets, comprise Numishare, which is likewise open-source and freely available through GitHub at <https://github.com/ewg118/numishare>. The seeds of Numishare were planted in 2007 with the digitisation and publication of the University of Virginia Art Museum Numismatic Collection, a project which was detailed in a paper presented at CAA 2009 in Williamsburg, Virginia. Numishare has evolved considerably since 2009, to the point where very few lines of code from the earliest era of the

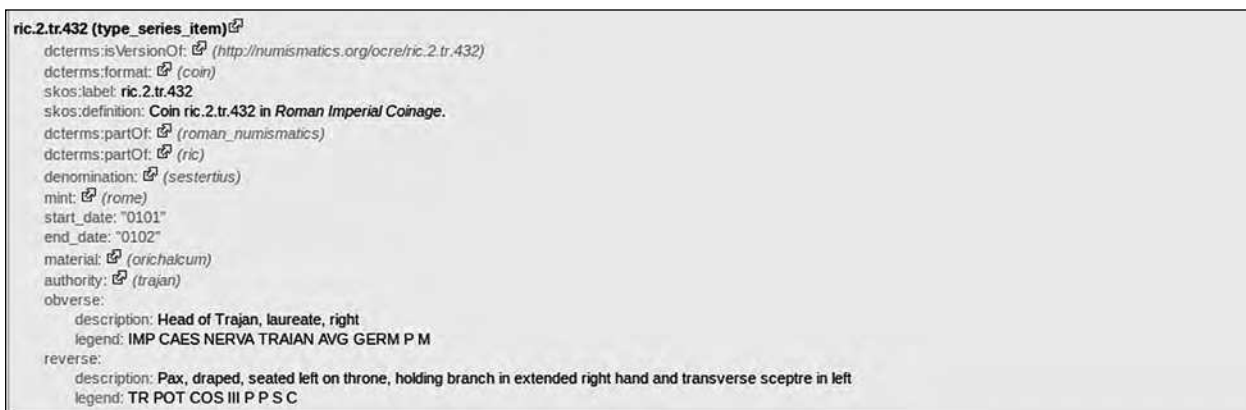


Figure 1. Nomisma coin type description.

project remain in the trunk of the repository. First and foremost, the XML data model was completely reinvented.

3.1 Numismatic Description Standard

The data model of the U. Va. Art Museum Numismatic Collection was an adaptation of Encoded Archival Description (EAD) to coins. EAD is an XML metadata standard used throughout the library and archival communities for encoding electronic finding aids.³ Thus, EAD is focused primarily on the description of written documents, though photographs, maps, and other predominantly non-textual objects can also be described. At the time, no standard for numismatic metadata had been authored, and the adaptation of EAD to coins was satisfactory given the needs of that project. Over time EAD was found to lack descriptors needed distinctly by numismatists and museum curators, as well as attributes to facilitate linked data interactions.

Independent of the development of EAD for coins and the U. Va. digitisation project, Sebastian Heath and Andrew Meadows of the American Numismatic Society had worked with a number of colleagues, at a series of workshops funded by the UK's AHRC, to develop a recommended list of numismatic and curatorial fields called the Numismatic Description (formerly Database) Standard (NUDS).⁴ NUDS contained many fields which were not included in EAD. Some of these were related to provenience and collection management: auction history, sale lot and price, owner, etc. Others were related to description of coins themselves: artists and engravers, symbols and countermarks, edge descriptions, and findspots. Many of these fields, like artist and engraver, can be expressed generally about the coin as a whole or more specifically about the obverse or reverse of the coin. NUDS fields were proposed with a database management system (DBMS) in mind, thus hierarchical complexities and repetition are difficult to represent easily or efficiently. Moreover, representing uncertainty

³ See "About EAD," The Library of Congress, last modified July 11, 2012, <http://www.loc.gov/ead/eadabout.html> for more information about the standard. The U.Va. collection is available at <http://coins.lib.virginia.edu>

⁴ Numismatic database fields detailed on Nomisma.org, accessed October 31, 2012, http://nomisma.org/nuds/numismatic_database_standard

on any level within the record is also difficult. The EAD and NUDS approaches had distinct advantages and disadvantages, and the American Numismatic Society has endeavored to address the differences between both in order to come to a solution that meets the requirements for all implementations of the model.

In summer 2011, work began on a data model which uses NUDS fields reconceived in an XML schema purpose-built for numismatics. By June 2012, the model, although still considered to be a draft, is stable, and several projects have been built on top of it, including OCRE.⁵ NUDS remains the label for the XML adaption of these fields, though the acronym now stands for "Numismatic Description Standard." The model influenced by other common XML standards found in the Library, Archive, and Museum world, like EAD, Encoded Archival Context – Corporate, Personal, and Family (EAC-CPF), VRA Core, Metadata Encoding and Transmission Standard (METS), Metadata Object Description Schema (MODS), and Text Encoding Initiative (TEI). Like several of these standards, NUDS uses the World Wide Web Consortium (W3C) xlink attributes for semantic linking, enabling links to URIs of Nomisma concepts, places on Geonames.org, and personal and corporate names defined by Virtual International Authority File (VIAF). A NUDS/XML record is fundamentally "linked data" because it is an electronic record which links to other resources on the web, although it is not linked data in the way that many of us may conceive of it: RDF in a triple store with a SPARQL endpoint for querying. Like traditional linked data architectures, querying for machine-readable data is supported in Numishare and will be discussed at a further point in this paper.

3.2 NUDS/XML: A synopsis

The document root of a NUDS/XML document is <nuds>, defined by the namespace <http://nomisma.org/nuds>. It requires one vital attribute for denoting the type of object represented by the electronic document. Objects are either physical or conceptual. A record for a physical object is self-explanatory: the XML document describes coin,

⁵ NUDS/XML data model draft detailed at <http://wiki.numismatics.org/nuds:nuds>, last edited October, 2012.

medal, token, paper note, or other numismatic object which physically exists in a collection. A conceptual object, however, is an intellectual construct. In numismatic terms, it is a coin type. Since OCRE is a corpus of coin types, this paper will focus primarily on describing the aspects of NUDS which define conceptual records.

Regardless of the physical or conceptual nature of the object being described, all NUDS documents must contain a NUDS Header. Like headers in TEI and EAD, the NUDS Header contains metadata about the electronic record itself: unique identifier, publication information, rights statements, and revision history. Within the document root, below the NUDS Header is Descriptive Metadata about the object itself. This includes, title, subjects, and physical, typological, undertype, findspot, and reference descriptions, in addition to administrative history for encoding provenience, ownership and accession information, and other other fields necessary for collections management. A physical object may utilize any or all of these categories of descriptors, but a coin type record would not contain descriptors for collections management, physical state, nor undertypes, which are only specific to physical objects. Thus, the Typological Description node is the only required Descriptive Metadata section.

The Typological Description (<typeDesc>) section of the NUDS document contains a mix of elements which utilize the W3C xlink:role and xlink:href attributes for semantic linking and other elements which simply contain free text. In Figure 2, the XML fragment shows that the <typeDesc> contains an object type, denomination, manufacture method, and material, each which link to concepts defined by Nomisma URIs. Below these elements are structured data: obverse, reverse, geographic, and authority nodes. The obverse and reverse may contain identical elements, including legend, type, and personal or corporate names defined by xlink:role. Roles may be authority, issuer, mint, region, deity, portrait, or others, as defined in Nomisma. The type element nested within obverse and reverse may contain repeatable description elements, with language defined by the xml:lang attribute. Numishare supports rendering documents in multiple languages, as long as labels denoted

by xml:lang exist either within the NUDS/XML record itself or within the RDF representation of the Nomisma URI.

As a practical example, suppose that a NUDS record contains both German and English type descriptions. A language parameter for German (“de”), may be passed to Numishare to display the appropriate type description, and lookups can be performed upon all Nomisma URIs within the document, parsing the RDF for each Nomisma ID and extracting the German SKOS preferred label for serialization into HTML. This move toward internationalization, greatly aided by Nomisma, will make it easier to develop and maintain multilingual interfaces for a single collection, as well as make it possible to query across multiple collections, regardless of the native language of the XML documents. When a NUDS record is published in Numishare and indexed into Solr, the URIs for Nomisma, the Pleiades Gazetteer of Ancient Places, Geonames, and VIAF IDs are stored in the index, enabling queries of the URIs directly (e. g., query all silver coins by the ID <http://nomisma.org/id/ar>, regardless of local encoding practice for materials). Furthermore, indexing of URIs for Pleiades places has facilitated integration with the PELAGIOS (Pelagios: Enable Linked Ancient Geodata In Open Systems) project, allowing Roman Imperial coins from Ephesus to be cross-searchable with other archaeological finds from the ancient city or classical texts which reference it⁶.

Below the Descriptive Metadata section within the NUDS document is the section for Digital Representations. In physical record types, one can embed a fragment of METS to link to digital images of a coin. Within OCRE, on the other hand, references to coins of that particular coin type can be inserted. The American Numismatic Society may have several physical specimens of a particular coin type, and other private or museum collections may also have specimens. Linking to physical coins which are represented online as stable URIs enables OCRE to ingest measurement data, findspots, and

⁶ In summation, PELAGIOS is a triplestore (and associated APIs) which aggregates RDF of various ancient collections online that are connected together via references to ancient places defined by Pleiades URIs. Read more at “About PELAGIOS,” The PELAGIOS Project, accessed October 31, 2012, <http://pelagios-project.blogspot.com/p/about-pelagios.html>.

```

<typeDesc>
  <objectType xlink:href="http://nomisma.org/id/coin">Coin</objectType>
  <date normal="-0018">18 BC</date>
  <denomination xlink:href="http://nomisma.org/id/denarius">Denarius</denomination>
  <material xlink:href="http://nomisma.org/id/ar">Silver</material>
  <obverse>
    <legend>CAESARI AVGVSTO</legend>
    <type>
      <description xml:lang="en">Head of Augustus, laureate, right</description>
    </type>
    <persname xlink:role="portrait">Augustus</persname>
  </obverse>
  <reverse>
    <legend>MAR VLT</legend>
    <type>
      <description xml:lang="en">Hexastyle domed temple, aquila between two standards inside</description>
    </type>
    <persname xlink:role="deity">Mars</persname>
  </reverse>
  <geographic>
    <geoname xlink:role="mint" xlink:href="http://nomisma.org/id/colonia_patricia">Colonia Patricia</geoname>
    <geoname xlink:role="region" xlink:href="http://nomisma.org/id/spain">Spain</geoname>
  </geographic>
  <authority>
    <persname xlink:role="authority" xlink:href="http://nomisma.org/id/augustus">Augustus</persname>
  </authority>
</typeDesc>

```

Figure 2. Typological Description fragment.

URLs to images, if available. These data can be used for quantitative and geographic analyses.

3.3 Numishare

Having introduced the NUDS/XML data model, it is now possible to discuss the applications which compose Numishare—how they function and how a user interacts with the public interface. As mentioned previously, there are four open-source, Java-based applications that run in Apache Tomcat. Essentially, Numishare is the code which links these separate applications together into one consolidated software suite. XSLT is the predominant scripting language within Numishare, with Javascript and CSS used to enhance the user experience. Apache Cocoon is the primary web publishing platform upon which the Numishare public interface is built. The Cocoon sitemap defines a myriad of pipelines which represent the Model-View-Controller relationship in XML applications development. Output serialization is formed by the combination of the data model (in some cases, NUDS/XML) and a controller and/or view (typically an XSLT stylesheet). The output serialization of a NUDS/XML document can be HTML rendered to a user of the OCRE interface in his or her browser, but XSLT stylesheets can also generate KML for rendering in maps or simply output the XML document itself to a user interested in taking advantage of the raw, unserialized data. The NUDS documents are stored in an eXist XML database, and therefore Cocoon reads the model directly through eXist’s REST interface.

Cocoon also handles interactions between the user and the Solr search index. According to the official website for Solr:

*“its major features include powerful full-text search, hit highlighting, faceted search, dynamic clustering, database integration, rich document (e.g., Word, PDF) handling, and geospatial search. Solr is highly scalable, providing distributed search and index replication, and it powers the search and navigation features of many of the world’s largest internet sites.”*⁷

Netflix was among its earliest adopters in private industry, and it has since become the *de facto* standard search index throughout the library community in the United States. It serves as the index for Blacklight, VuFind, and various Fedora repository implementations, as well as countless other projects. Solr has been a part of Numishare’s application stack since 2007. It was chosen as the search index as a means of improving the user experience in searching and browsing that were seen as shortcomings of Numishare’s two main influences, the Princeton University Numismatic Collection and the Berlin Münzkabinett, which were in 2007 and continue to be among the most visible numismatic databases on the web.⁸

7 “Apache Solr,” Apache Software Foundation, accessed June 20, 2012, <http://lucene.apache.org/solr/>.

8 See “RBSC: Numismatic Collection,” Princeton University, accessed October 31, 2012, <http://www.princeton.edu/~rbsc/department/numismatics/> and “Münzkabinett der Staatlichen Museen Berlin,” accessed October 31, 2012, <http://www.smb.museum/ikmk/>.

The following is a brief description of the functionality of the search and browse interfaces in Numishare. Solr accepts queries following the Lucene query syntax submitted by the user through the interface, and Solr responds through its REST API with query results in the form of XML (although it also supports JSON output). Cocoon intertwines this XML model with an XSLT stylesheet, which renders the results in the form of human-readable HTML. Query results can also be serialized into several other formats. Results can be delivered in the form of an Atom feed, which allows users to subscribe to changes made to coins which fit their key area of interest (e. g., coins of Hadrian). More importantly, the Atom feed is the primary means by which data may be harvested from OCRE by programmatic means. The feed contains 100 hits per page, with links to next and previous pages, as well as links to alternates to the default HTML serialization of a coin type, such as RDF, KML, and NUDS/XML. Therefore, it is possible for a script to page through the entire collection and extract the machine-readable NUDS metadata. In addition to the Solr-based Atom feeds, Numishare also supports serializing Solr search results as KML and CSV. While Cocoon handles output for the public user interface, XForms applications processed by Orbeon, which also runs in Tomcat, form the foundation of the private administrative interface.

XForms is a W3C specification for next-generation web forms which adhere to the Model-View-Controller philosophy.⁹ Orbeon is used in the administrative back-end to edit complex XML metadata and interact with REST services. XForms applications query various APIs on the web for enhancing controlled vocabulary and manage Create, Read, Update, and Delete (CRUD) operations between the eXist database and Solr index. Specific details about the architecture of Numishare's back-end are too complex to be detailed within the limited space of this paper, but advantages are briefly enumerated below:

- XForms supports validation. For example, numeric fields such as diameter and weight must adhere to the `xs:decimal` datatype. In other cases, XForms validation warns the user that certain

⁹ The XForms Wikipedia article provides a sufficient description of the standard. "XForms," Wikipedia.org, last modified October 21, 2012, <http://en.wikipedia.org/wiki/XForms>.

fields may not be left blank. Invalid data cannot be sent into eXist.

- Authority control. XForms taps into auto-suggest for local controlled vocabulary delivered from the Solr index or interacts with REST APIs of Geonames, VIAF, and Nomisma.org to embed URIs directly into the `xlink:href` attribute within the NUDS/XML document.
- Link to images with Flickr's APIs.
- Easily link to coins in other collections through automated processing.
- A fuller description of XForms, as well as other use cases for similar applications are detailed in "XForms for Libraries, An Introduction," published in *Code4Lib* 11.¹⁰

Together, this suite of applications known as Numishare forms the basis for a powerful research tool for those interested in the study of Roman Imperial Coins. Several features of the public interface have already been discussed (KML and Atom feeds), but this paper will now turn toward a more focused discussion of the public interface of OCRE.

4. OCRE: A Powerful Research Tool for Roman Imperial Coinage

4.1 Searching and browsing

Despite numerous and significant advances in Numishare since its inception, especially with respect to collections management and publishing in the XForms-driven backend, the application in its core is designed to deliver a simple, yet sophisticated interface for navigating and interacting with coin collections. The search and browse interfaces are good places to begin a discussion about the features of Numishare, for these are what separates Numishare from most other collections. Solr supports faceted searching based on a handful of typological categories: authority, deity, denomination, manufacture, issuer, material, mint, portrait, and region. Full-text search is supported,

¹⁰ Ethan Gruber, Chris Fitzpatrick, Bill Parod, and Scott Prater, "XForms for Libraries, An Introduction," *Code4Lib* 11 (September 21, 2010), accessed October 31, 2012, <http://journal.code4lib.org/articles/3916>.



Figure 3. Preview of dynamic mapping interface.

as well as searches of other fields, like obverse or reverse legend or type, date ranges, among others. Search results can be sorted by a handful of fields, including all those listed above.

Importantly, Solr supports wild card searches, which are especially useful legends which are only partially legible. Such searches can aid in the identification of coins acquired in archaeological excavations which may be in a poor state of preservation. For example, one may search for “P?CI AVG*” within the reverse legend to yield several dozen matches for the legend “PACI AVGVST” or “PACI AVGVSTAE.” A “?” wild card represents zero or one unknown characters while “*” represents one or more. Results may be filtered further by facets. Given this partial legend query, suppose the metal of bronze is readily identifiable, and the obverse portrait, however worn, appears to exhibit the features of the emperor, Domitian. In OCRE, this yields one result (through the reign of Antoninus Pius, ending in A.D. 161). Such queries can be represented by Atom, and the metadata for coin types matching these criteria can be ingested into local databases by machine-readable processes, reducing the data entry workload of archaeologists or other recorders.

With the adaptation of Numishare to the American Numismatic Society’s collection, called Mantis, in early 2011, a new faceted geographic search interface was introduced, in addition to the traditional textual interface typically associated with Solr searches. Javascript was written to connect

faceted search with OpenLayers, an open-source Javascript library for mapping. These Javascript functions accept interactions from the user of the interface to query Solr, which replies with XML piped through a Cocoon+XSLT transformation into KML to rapidly update the OpenLayers map. Visualizing queries geographically is potentially immensely useful. One may map the distribution of coins over time and space (Figure 3 shows the mints which produced coins depicting the personification, Pax) or map the distribution of a

particular coin type over the whole of the Roman Empire, which would give an economic historian a glimpse of ancient trade networks.

The final interface constructed upon Solr is one that visualizes search results in the form of charts and graphs. A query of all coins which depict the deity, Victory, may be visualized as a bar graph depicting the total counts per Roman emperor. In this, Vespasian leads all other emperors.¹¹ Such visualisations may lead to research questions which may have otherwise never been asked.

4.2 Coin type records

Like most other aspects of Numishare, the coin or coin type record HTML serialization has evolved considerably since it was initially documented at CAA 2009, although some features of the HTML page have remained. The design of the object record page was conceived in 2007 to address the perceived shortcomings of record pages in other databases, namely the “dead-end” nature of those pages. Links from search results typically led to pages for coins which contained metadata, but no way to seamlessly traverse from those records to records of related coins. The 2007 design created links for typological attributes—like denomination, material, and mint—which would direct the user to the search results page for that attribute. One could navigate from a particular denarius of Augustus to other denarii

¹¹ Visualisation parameters are passed RESTfully to the URL. See http://numismatics.org/ocre/visualize?limit=32&type=areaspline&category=authority_facet&q=deity_facet%3A%22Victory%22 for a representation of this query.

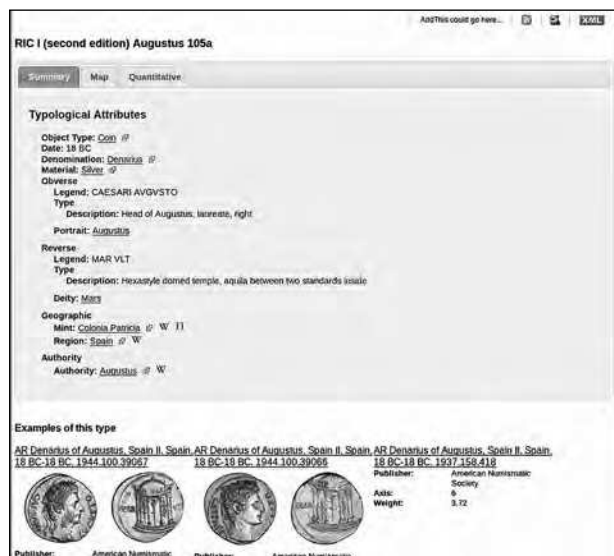


Figure 4. A coin type and its associated coins.

or other coins of Augustus. OCRE continues to provide this functionality, which had unintended, but nevertheless positive, consequences. Google and other robots are able to crawl from record to record through RESTful search results to make records available through search engines.

To enhance the user experience, when Cocoon builds the HTML representation of a NUDS/XML document for a coin type (Fig. 4), lookups are performed upon each Nomisma URI within the XML to read the RDF data stream for the given ID. Each related web resource (designated by a `skos:related @rdf:resource`) listed in the RDF is transformed into a link within the coin type HTML page, enabling the user to visit those resources to gather more contextual information about the Nomisma concept (e. g., Wikipedia articles). Additional features of the HTML page for a coin type includes links to associated physical objects in other collections, an OpenLayers map which renders points for mints and known findspots associated with the coin type, and a section for quantitative analysis of the type. While fairly rudimentary at the present stage, especially given the limited number of physical coins associated with each coin type, this section of the page lists the average weight for the coin type and allows the user to create graphs which compare the average weight of the specified coin type with the average weights of coins with similar attributes (e. g., the average weight of other bronze

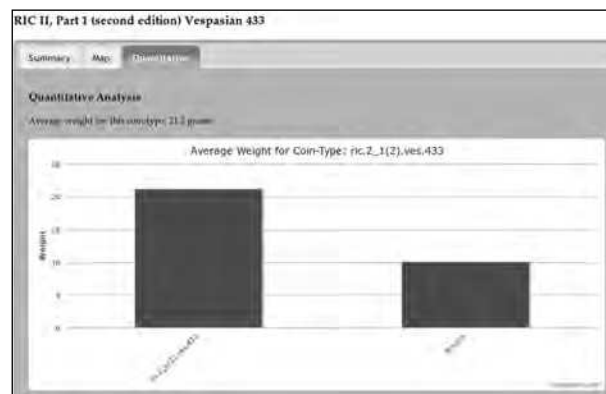


Figure 5. Weight comparison.

Roman coin-types compared to a bronze sestertius of Vespasian, no. 433: Fig. 5).¹² This interface for visualizing weights will be expanded to encompass other measurements, as well as standard deviations.

Lastly, the record page links to other serializations of the record: NUDS/XML, Atom, RDF, and KML.¹³ These links are established with icons represented visually in the browser, as well as in the HTML header to facilitate machine reading.

5. Conclusions

OCRE represents a fundamental shift in Roman numismatics. OCRE will be freely available to everyone with Internet access, creating a low-cost and portable tool accessible whether in the field or in a library. The public user interface serves as a powerful tool for searching, sorting, and visualizing Roman coinage in a manner that is simply impossible in the printed catalogues. Since the coin type metadata is open and can be considered to be canonical representations of entries in RIC, other collections may make use of OCRE's APIs for extracting this metadata for ingestion or reference in their own databases, saving time in the data entry process and deferring the burden of maintaining the data to the OCRE project editors. In theory, OCRE can be used by field archaeologists to quickly identify excavated coins, aiding in the classification

¹² Parameters for analyses are passed RESTfully into the URL. Go to http://numismatics.org/ocre/id/ric.2_1%28%29.ves.433?type=column&weightQuery=material_facet%3A%22Bronze%22#charts to see this chart in the browser.

¹³ The RDF ontology conforms to the PELAGIOS specification, but CIDOC-CRM conforming to The British Museum template will likely be implemented within the next twelve months.

and interpretation of stratigraphic contexts, even when the excavation lacks numismatic specialists.

OCRE will continue to evolve even after the publication of this paper. More NUDS records will be added into OCRE, terminating with Anastasius in the early sixth century. New features for quantitative analyses and visualisations will be added into Numishare. Physical coins in other collections will continue to be linked to coin types. At this moment, the Portable Antiquities Scheme, British Museum, OpenContext and the Münzkabinett of the Staatlichen Museen in Berlin have committed to sharing their data with the project. Since a relative small proportion of the American Numismatic Society's coins linked to OCRE have attested findspots, it is a primary goal moving forward to import data about archaeologically excavated coins into the project. Only then can the full scholarly potential of OCRE be realised. In some sense, OCRE, at the date of this publication, is merely a starting point in the development of one of the most comprehensive and sophisticated tools for the study of Roman coins on the web.

References

- Apache Software Foundation. 2012. "Apache Solr." Accessed June 20, 2012. <http://lucene.apache.org/solr/>.
- Gruber, E. 2009. "Encoded Archival Description for Numismatic Collections." Paper presented at the *Computer Applications and Quantitative Methods in Archaeology Conference*, Williamsburg, Virginia 23-27 March, 2009.
- Gruber, E. 2012. "Numismatic Description Standard: Draft." Last modified October, 2012. <http://wiki.numismatics.org/nuds:nuds>.
- Gruber, E, C. Fitzpatrick, B. Parod, and S. Prater. 2010. "XForms for Libraries, An Introduction." *Code4Lib* 11. <http://journal.code4lib.org/articles/3916>.
- The Library of Congress. 2012. "About EAD." Last modified July 11, 2012. <http://www.loc.gov/ead/eadabout.html>.
- Linkeddata.org. 2012. "Linked Data - Connect Distributed Data across the Web." Accessed June 20, 2012. <http://linkeddata.org/>.
- Nomisma.org. 2012. "Numismatic Description Standard." Accessed June 20, 2012. http://nomisma.org/nuds/numismatic_database_standard.
- PELAGIOS Project. 2012. "About PELAGIOS." Accessed October 31, 2012. <http://pelagios-project.blogspot.com/p/about-pelagios.html>.
- Wikipedia.org. 2012. "XForms." Last modified October 21, 2012. <http://en.wikipedia.org/wiki/XForms>.
- World Wide Web Consortium. 2012. "XML Linking Language (XLink) Version 1.1." Accessed June 20, 2012. <http://www.w3.org/TR/xlink11/>.