

Connecting the Smithsonian American Art Museum to the Linked Data Cloud*

Pedro Szekely¹, Craig A. Knoblock¹, Fengyu Yang², Xuming Zhu¹,
Eleanor E. Fink¹, Rachel Allen³, and Georgina Goodlander³

¹ University of Southern California

Information Sciences Institute and Department of Computer Science
pszekely@isi.edu, knoblock@isi.edu, xumingzh@usc.edu, efink@ifc.org

² School of Software, Nanchang Hangkong University, China
frueyang@gmail.com

³ Smithsonian American Art Museum
rallen@si.edu, goodlanderg@si.edu

Abstract. Museums around the world have built databases with metadata about millions of objects, their history, the people who created them, and the entities they represent. This data is stored in proprietary databases and is not readily available for use. Recently, museums embraced the Semantic Web as a means to make this data available to the world, but the experience so far shows that publishing museum data to the Linked Data cloud is difficult: the databases are large and complex, the information is richly structured and varies from museum to museum, and it is difficult to link the data to other datasets. This paper describes the lessons learned in publishing the data from the Smithsonian American Art Museum (SAAM). We highlight complexities of the database-to-RDF mapping process, discuss our experience linking the SAAM dataset to hub datasets such as DBpedia, and present our experience in allowing SAAM personnel to review the information to verify that it meets the high standards of the Smithsonian. Using our tools, we helped SAAM publish 5-star Linked Data of their complete holdings (41,000 objects, 8,000 artists), richly linked to DBpedia and the Getty Union List of Artist Names (ULAN), and verified to be of high quality.

1 Introduction

Recently, there have been a number of efforts to publish metadata about the objects in museums as Linked Open Data (LOD). Some notable efforts include the Europeana project [7], which published data on 1,500 of Europe's museums,

* This research was funded by the Smithsonian American Art Museum. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Smithsonian Institution.

libraries, and archives, the Amsterdam Museum[3], which published data on 73,000 objects, and the LODAC Museum [11], which published data from 114 museums in Japan. Despite the many recent efforts, there are still significant challenges in publishing data about artwork to the Linked Data Cloud.

There are three steps in the process of mapping the data of a museum to the Linked Data Cloud. First, there is the problem of mapping the underlying data sources with the metadata about the artwork into RDF. As noted, by de Boer et al. [3], this can be a complicated process since many museums have rich, hierarchical, graph-structured data that has been collected over many years. There are often attributes or properties of this data that are unique to a particular museum and the data is often inconsistent and noisy since it has typically been maintained over a long period of time by many individuals. In past work, this mapping process is typically done by manually writing rules or programs to define the mapping. Second, once the data is in RDF, the next step is to find the links from the metadata to other repositories, such as Dbpedia or Geonames. In previous work, this has also been done by defining a set of rules for performing the mapping. Because of the difficulty of this problem, the number of links in past work in this space is actually quite small as a percentage of the total set of objects that have been published. The third step is to curate the data to ensure that both the information that is published and the links to other sources within the LOD are accurate. Because curation is so labor intensive, this step has been largely ignored in previous work.

The general goal of our project is to develop the technology to allow museums and other organizations to map their own data to Linked Open Data. The work described in this paper is the first step towards that larger goal. In particular, we describe the lessons learned in the process of mapping the metadata that describes the 41,000 objects of the Smithsonian American Art Museum (SAAM). In our previous work, we developed a system called KARMA for mapping complex, hierarchically structured sources to RDF with respect to a given domain ontology. However, in the real-world data provided by the Smithsonian, we discovered that there were complex structures that required new capabilities in KARMA. In terms of linking, we found that mapping the entities, such as artist names, to DBpedia could not be easily or accurately performed using existing tools, so we developed a specialized mapping approach to achieve high precision matches. In the past work in this space, either very few of the entities are linked across sources or when they are linked the precision is quite low. Finally, to ensure that the Smithsonian publishes high quality linked data, we developed a curation tool that allows the Museum staff to easily review and correct any errors in the automatically generated links to other sources.

In the remainder of this paper, we describe our approach and present the lessons learned in mapping (Section 2), linking (Section 3), and curating (Section 4) the data of the Smithsonian American Art Museum. For each of these topics, we describe our approach, present lessons learned, and evaluate the effectiveness of our approach. We then compare our work to previous work (Section 5) and conclude with a discussion of the contributions and future work (Section 6).

2 Mapping the SAAM Data to RDF

2.1 The SAAM Database

SAAM stores collection metadata in a relational database managed by TMS⁴, a comprehensive data management system for museums. The SAAM deployment of TMS consists of over 100 tables, containing significant amounts of data that needs to remain private. In order to avoid issues about private data, we only use the tables that the museum uses to drive their Web site. All the information in these eight tables already appears on the museum Web site, so the museum is amenable to publishing it as Linked Data. The structure and format of these data are tailored to the needs of the Web site and some fields need to be decoded to produce appropriate RDF. For example, descriptive terms are encoded in text such as “Authorities\Attributes\Objects\Folk Art”. The database includes data about 41,267 objects and the 8,261 artists who created them.

For objects, the database contains the regular tombstone information including classification (e.g., sculpture, miniature), their role (e.g., artist, printer), available images, terms (e.g., Portrait Female – Turner, Tina). For artists, the database contains names, including multiple variants (e.g., married name, birth or maiden name), title and suffixes, biographical information and geographical information including city, county, state and country of relevant places (e.g., birth and death place, last known residence) and citation information.

Lesson 1: Satisfy the legal department first. Much of the data in museums is proprietary and getting approval from the legal department can be challenging. We use the data that drives the Web site; it is not the raw data, but adequate and circumvents issues that could have stopped or delayed the project.

2.2 Europeana Data Model (EDM)

The Europeana Data Model (EDM) is the metamodel used in the Europeana project⁵ to represent data from Europe’s cultural heritage institutions. We base our ontology (Figure 1) on EDM because it is a comprehensive, yet extensible OWL ontology and it maximizes the opportunity to integrate the SAAM Linked Data with other datasets that are already using it. In addition to EDM, we use other popular vocabularies/ontologies used in the Linked Data cloud: *SKOS*⁶ for the classification of artworks, artist and place names; *Dublin Core*⁷ for the tombstone data; *RDA Group 2 Elements*⁸ to represent biographical information; and *schema.org*⁹ to represent geographical data (city, state, country).

⁴ gallerysystems.com/tms

⁵ <http://europeana.eu>

⁶ <http://www.w3.org/2004/02/skos/>

⁷ <http://purl.org/dc/elements/1.1/> and <http://purl.org/dc/terms/>

⁸ <http://rdvocab.info/ElementsGr2>

⁹ <http://schema.org/>

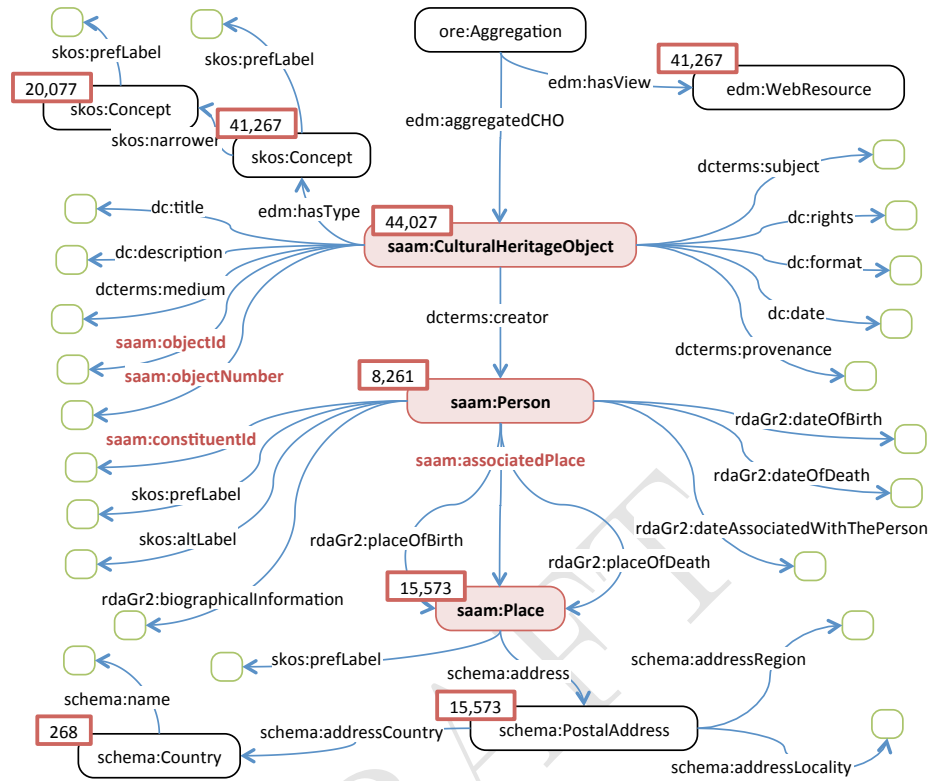


Fig. 1. The SAAM ontology. Named ovals represent classes, un-named green ovals represent literals, arcs represent properties, boxes contain the number of instances generated in the SAAM dataset.

One of the most challenging tasks in the project was selecting and extending the ontologies. We considered EDM and CIDOC CRM¹⁰; both are large and complex ontologies, but neither fully covers the data that we need to publish. We needed vocabularies to represent biographical and geographical information, and there are many to choose from. Following the lead of the Amsterdam Museum [3], we used RDA Group 2 Elements for the biographical information. We didn't find guidance for representing the geographical information in the cultural heritage community so we selected schema.org as it is a widely used vocabulary. Our extensions (shown in boldface/red in Figure 1) are subclasses or subproperties of entities in the ontologies we reuse.

Lesson 2: A library of ontologies for cultural heritage is desperately needed. While EDM represents an excellent starting point for modeling cultural heritage data, the community can benefit from guidance on vocabularies to represent data not covered by EDM and an integrated library with the recommended ontologies.

¹⁰ <http://www.cidoc-crm.org>

2.3 Using KARMA to Map the SAAM Data to RDF

In previous work [9], we developed KARMA, a tool to map structured data to RDF according to an ontology of the user’s choice. The goal is to enable data-savvy users (e.g., spreadsheet users) to do the mapping, shielding them from the complexities of the underlying technologies (SQL, SPARQL, graph patterns, XSLT, XPath, etc). KARMA addresses this goal by automating significant parts of the process, by providing a visual interface (Figures 2 to 4) where users see the KARMA-proposed mappings and can adjust them if necessary, and by enabling users to work with example data rather than just schemas and ontologies.

The KARMA approach to map data to ontologies involves two interleaved steps: one, assignment of *semantic types* to data columns and two, specification of the relationships between the semantic types. A semantic type can be either an OWL class or the range of a data property (which we represent by the pair consisting of a data property and its domain). KARMA uses a conditional random field (CRF) [10] model to learn the assignment of semantic types to columns of data [5] from user-provided assignments. KARMA uses the CRF model to automatically suggest semantic types for unassigned data columns. When the desired semantic type is not among the suggested types, users can browse the ontology through a user friendly interface to find the appropriate type. KARMA automatically re-trains the CRF model after these manual assignments.

The relationships between semantic types are specified using paths of object properties. Given the ontologies and the assigned semantic types, KARMA creates a graph that defines the space of all possible mappings between the data source and the ontologies [9]. The nodes in this graph represent classes in the ontology, and the edges represent properties. KARMA then computes the minimal tree that connects all the semantic types, as this tree corresponds to the most concise model that relates all the columns in a data source, and it is a good starting point for refining the model. Sometimes, multiple minimal trees exist, or the correct interpretation of the data is defined by a non-minimal tree. For these cases, KARMA provides an easy-to-use GUI to let users select a desired relationship (an edge in the graph). KARMA then computes a new minimal tree that incorporates the user-specified relationships.

Challenge 1: Data preparation. We encountered multiple situations where we had to filter and transform data prior to modeling it and converting it to RDF. The following are the the types of data preparation tasks we encountered: *Filtering tables:* for example, the SAAM tables represent constituents, which includes both people and organizations. The ontologies for people and organizations are different so we defined database views to filter the tables accordingly. *Data extraction:* for example, the keywords associated with the art objects need to be extracted from text such as “Authorities\Attributes\Objects\Subject Specific\Animal\bird\owl”. *Concatenating and formatting columns:* the SAAM tables represent people names, dates and places in a structured form (e.g., year, month and date in separate columns). We needed to concatenate these fields to

construct values for single properties (e.g., `dateOfBirth`), taking care to insert separators and leading zeroes to format them appropriately.

We addressed these data preparation tasks before modeling them in KARMA by writing scripts in Java to transform the data and by defining views in SQL. We then loaded the new tables and views in KARMA to model them. While data preparation is routine in database applications and powerful tools are available to support them, RDF mapping tools (including KARMA) lack the needed expressivity. Tools like ClioPatria [3] allow users to define expressions in a full programming language (Prolog in the case of ClioPatria) and invoking them within their mapping rules. Our approach is to enable users to use whatever tools they are familiar with in a prior data preparation step.

Lesson 3: The data preparation/data mapping split is effective: The range of data preparation tasks is open-ended and ad hoc. It is wise to acknowledge this and to design a data mapping architecture that is compatible with traditional data preparation tools. This allows the data mapping language to remain relatively simple. KARMA integrates with a data preparation step by providing the ability to specify many aspects of the mapping in the data tables themselves (discussed below). We did the data preparation primarily using SQL views, other users of KARMA have used Google Refine¹¹.

Challenge 2: Mapping columns to classes. Mapping columns to the ontology is challenging because in the complete SAAM ontology there are 407 classes and 105 data properties to choose from. KARMA addresses this problem by learning the assignment of semantic types to columns. Figure 2 shows how users define the semantic types for the `constituentid` (people or organizations) and `place` columns in one of the SAAM tables. The figure shows a situation where KARMA had learned many semantic types. The left part shows the suggestions for `constituentid`. The SAAM database uses sequential numbers to identify both constituents and objects. This makes them indistinguishable, so KARMA offers both as suggestions, and does not offer other irrelevant and incorrect suggestions. The second example illustrates the suggestions for the `place` column and shows how users can edit the suggestions when they are incorrect.

¹¹ <http://code.google.com/p/google-refine/>

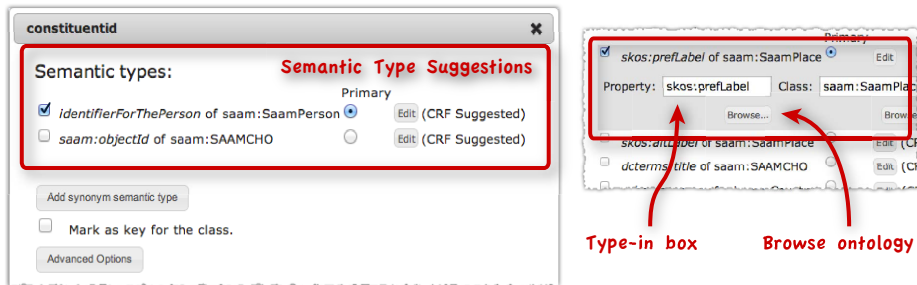


Fig. 2. KARMA suggests semantic types and enables users to easily customize them.

Challenge 3: Connecting the classes. This is also challenging because there are 229 object properties in the ontology to choose from. Figure 3 illustrates how KARMA automatically connects the semantic types for columns as users define them. In the first screen the user assigns a semantic type for `constituentid`. In the second screen, the user assigns a semantic type for `place`, and KARMA automatically adds to the model the `associatedPlace` object property to connect the newly added `SaamPlace` to the pre-existing `SaamPerson`. Similarly, when the user specifies the semantic type for column `city`, KARMA automatically adds the `address` object property.

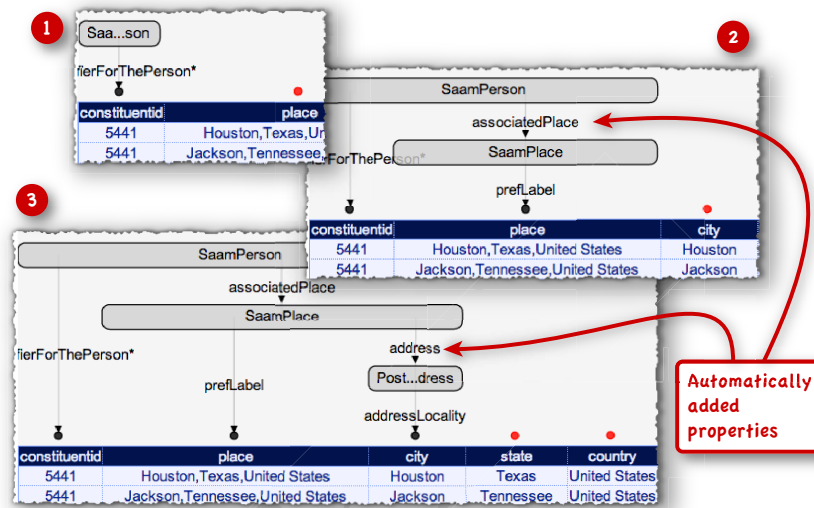


Fig. 3. Each time the user adds new semantic types to the model, KARMA connects them to the classes already in the model.

Each time users model the semantic type of a column, KARMA connects it to the rest of the model. In the examples, the connections use a single property, but KARMA searches the whole graph and finds longer paths when appropriate. In addition, weights in the graph [9] bias the algorithm to prefer specific properties rather than general properties inherited from superclasses. Sometimes, multiple lowest-cost models exist, or the appropriate model is not the lowest-cost model. Users can easily adjust the proposed model by clicking on an incorrect property and selecting the appropriate one from a menu of all compatible properties.

Lesson 4: Property domain and range definitions are important: KARMA leverages domains and ranges to automate modeling the relationships between the columns in the tables, often selecting the correct property. When KARMA proposed non-sensical, complicated paths to connect classes (e.g., subclass path via Thing), it was often because the relevant properties lacked the proper domain or range information or because the classes we defined had not been defined as subclasses of the appropriate classes. This feedback helped us to integrate the SAAM-specific ontology with the large complex ontologies we are reusing.

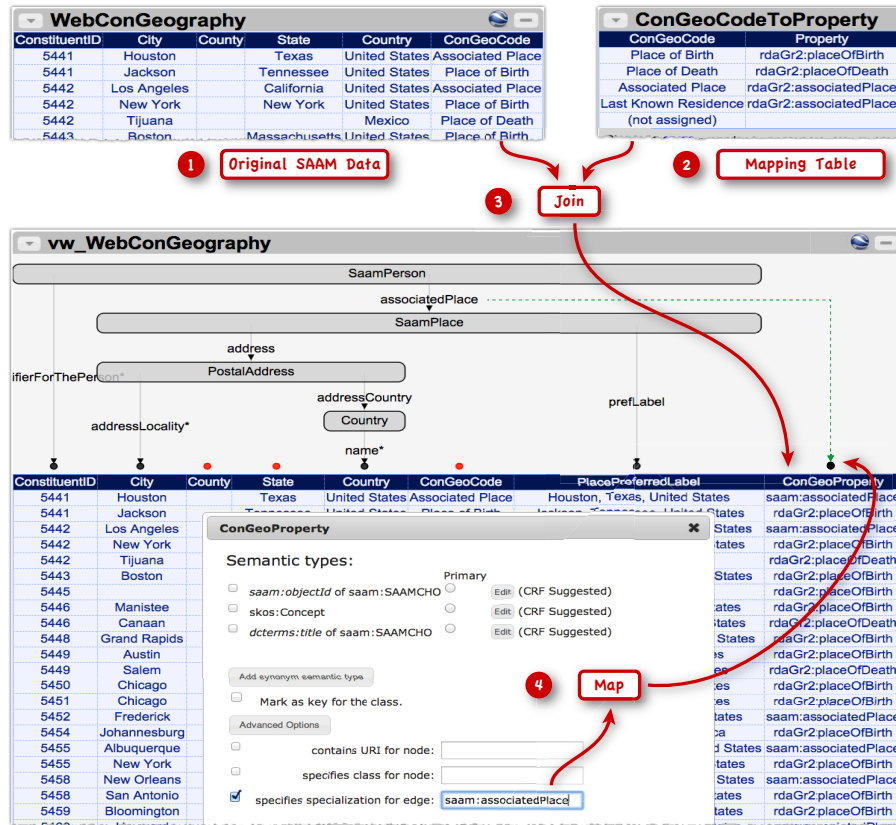


Fig. 4. Mapping columns where different rows must be mapped using different properties: 1) the data table; 2) a table to translate ConGeoCode to ontology properties; 3) join to add a column with the property URIs; 4) map values in the ConGeoProperty column to the associatedPlace property.

Challenge 4: Mapping depends on field values. Figure 4 illustrates a situation where the mapping from a table to the *desired* ontology cannot be specified at the schema level. The WebConGeography table contains information associated with people. Each row represents a place association: the first column (ConstituentID) represents the person identifier, the middle columns represent the place and the last column (ConGeoCode) represents the meaning of the place. The SAAM ontology defines a generic property associatedPlace to represent the relationship between a person and a place. This general property is appropriate for the first and third rows, but not for the others (e.g., the second row should use the more specific property rdaGr2:placeOfBirth).

To model situations such as this one, users add a column that contains the required data. In the particular case illustrated in Figure 4, the user can define a table that maps the ConGeoCodes to the appropriate properties (step 2) and then do a join to add the new column (step 3). Finally, when defining the semantic

type for the new column (step 4), users can specify that the values in the column specialize the `associatedPlace` property. Analogous situations arise in the SAAM tables that represent data about dates associated with people and variant people names. This type of mapping can be defined in tools such as D2RQ¹², but requires an expert user to define multiple, complex conditional mapping rules.

Lesson 5: Supporting row-level meta-data solves many complex mapping problems. The same mechanism we used to model row-specific properties can be used to model row-specific classes, URIs and language tags. It enables users to invoke arbitrary computation using their favorite tools to define data-dependent aspects of the mapping that cannot be cleanly represented in declarative representations. Other approaches such as D2RQ offer a limited set of built-in functions (e.g., concatenation, regular expression) that can be extended by writing Java classes. Our approach enables users to use whatever tools they are comfortable using.

Evaluation We evaluated the effectiveness of KARMA by mapping 8 tables (29 columns) to the SAAM ontology (Table 1). We performed the mapping twice: in *Run 1*, we started with no learned semantic types, and in *Run 2* we ran KARMA using the semantic types learned in the first run. The author of the paper that designed the ontology performed the evaluation. Even though he knows which properties and classes to use, when KARMA didn't suggest them he used the browse capability to find them in the ontology instead of typing them in. It took him 18 minutes to map all the tables to RDF, even in the first run, when KARMA's semantic type suggestions contained the correct semantic type 24% of the time. The second run shows that the time goes down sharply when users don't need to browse the ontology to find the appropriate properties and classes. The evaluation also shows that KARMA's algorithm for assigning relationships among classes is very effective (85% and 91% correct in *Run 1* and *Run 2*).

Lesson 6: Ontology design is the hard part. Even though it takes about 8 to 18 minutes to map all the tables using KARMA, it took about 2 weeks after the initial design of the ontology to map all the tables. We spent the time designing and redesigning the ontology. During that period, we mapped the tables many times to slightly different ontologies. So, in Table 1 *Run 2* is typical as we spent significant time rerunning KARMA after slight changes to the ontology.

¹² <http://d2rq.org>

Table 1. Effectiveness of KARMA's automation capabilities.

	<i># of times KARMA's top 4 suggestions contain the correct semantic type</i>	<i># of times KARMA correctly assigns relationships among classes</i>	<i>Time (minutes)</i>
<i>Run 1</i>	7 out of 29 (24%)	30 out of 35 (85%)	18
<i>Run 2</i>	27 out of 29 (93%)	32 out of 35 (91%)	8

3 Linking to External Resources

The RDF data will benefit the Smithsonian museum and the community if it is linked to useful datasets. We focused on linking SAAM artists to DBpedia¹³ as it provides a gateway to other Linked Data resources and it is a focus for innovative applications. We also linked the SAAM artists to the Getty Universal List of Artist Names (ULAN) and to the artists in the Rijksmuseum dataset.

SAAM needs to verify all matches before publishing them, so to make the verification process manageable, we sought high-precision algorithms. We matched people using their names, including variants, and their birth dates and death dates. The task is challenging because people’s names are recorded in many different ways, multiple people can have the same name, and birth dates and death dates are often missing or incorrect. For technical reasons we did not use other information, although we plan to do so in the future.

Our approach involves estimating the ratio of people in DBpedia having each possible value for the properties we use for matching (e.g., ratio of people born in 1879). For attributes compared using equality (birth/death years), we scan all people in DBpedia counting the number that have each specific value. For dependent attributes such as birth and death year, we also compute the ratios for pairs of values. We compare names using the Jaro-Winkler string metric [4], and for them compute the ratios as follows: we divide the interval $[0, 1]$ in bins of size ϵ , and for each bin we estimate the number of pairs of people whose names differ by a Jaro-Winkler score less than ϵ . Empirically, we determined that $\epsilon = 0.01$ and 10 million samples yield good results in our ground truth dataset.

The matching algorithm is simple. Given a SAAM person and a DBpedia person, their matching score is $s = 1 - d * n$ where d is the *date score* and n is the *name score*. If the dates match exactly, d is the fraction of people in DBpedia with those dates. If dates do not match exactly, d is the sum of the fractions for all the intervening years. n is the fraction of people in DBpedia whose Jaro-Winkler score is within ϵ from the score between the given pair of people, computed using the estimates discussed above. We use a simple blocking scheme based on the first letter of the name. To increase efficiency, we discard pairs whose birth or death years differ by more than 10 and whose Jaro-Winkler score is less than 0.8.

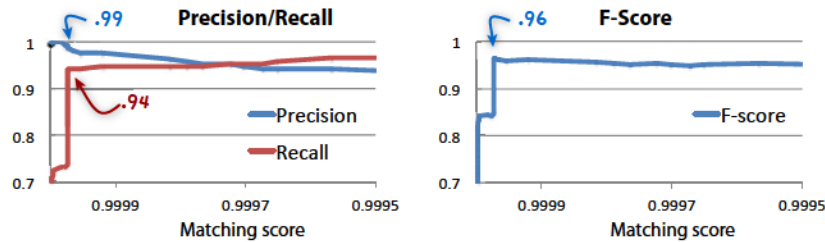


Fig. 5. Evaluation results.

¹³ dbpedia.org

Evaluation To evaluate our algorithm we constructed ground truth for a dataset of 535 people in the SAAM database (those whose name starts with A). We manually searched in Wikipedia using all variant names and verified the matches using the text of the article and all fields in the SAAM record, including the biography. We found 176 matches in DBpedia.

Figure 5 shows the evaluation results on the ground truth (note that the matching score s decreases from left to right). The highest F-score .96 achieves a precision of .99 and a recall of .94 (166 correct results, 1 incorrect result). At this threshold s^* all names and the years for all but 5 people match exactly. The incorrect result is one where neither date matches, but interestingly, there are 4 results where the years are present but not equal. The sharp increase in recall comes at a score $s > s^*$ where suddenly 37 results for people missing a single date are above threshold (all these are correct results). The next interesting threshold is $\hat{s} = 0.9995$. Between s^* and \hat{s} are 13 results; of these, 4 are correct (2 with non matching names) and 9 incorrect, yielding .938 precision and .966 recall. For $s < \hat{s}$, the density of correct results decreases sharply, containing only 4 correct results in the next 286 candidates. Based on these results, we used \hat{s} as the threshold to match the SAAM data against all DBpedia people (2,807 results), the Getty ULAN (1,759 results) and the Rijksmuseum (321 results).

4 Curating the RDF and the Links

Museums need the ability to ensure that the Link Data they publish are of high quality and that quality remains high as data sources evolve.

The first aspect of the curation process is to ensure that the RDF is correct. Museum personnel can easily browse individual RDF records on the Web, but without seeing and understanding the relationship between an RDF record and the underlying database records, it is hard to assess whether the RDF is correct. KARMA helps museum personnel understand these relationships at the schema level by graphically showing how database columns map to classes and properties in the ontology (e.g., Figures 3 and 4). KARMA also lets users click on individual worksheet cells to inspect the RDF generated for it, helping them understand the relationships at the data level. These graphical views also enabled SAAM personnel and the Semantic Web researchers to communicate effectively while refining the ontology and the mappings. Our goal by the end of the project is that SAAM personnel will use KARMA to refine the mappings on their own.

The second aspect of the curation process is to ensure that links to external sources are correct. Our approach is to 1) record the full provenance of each link so that users (and machines) can record links and inspect them when the data sources or the algorithm change, and 2) make it easy for users to review the results of the linking algorithm. We use the PROV ontology¹⁴ to represent provenance data for every link including revisions, matching scores, creation times, author (human or system/version) and data used to produce the link.

¹⁴ <http://www.w3.org/TR/prov-o/>

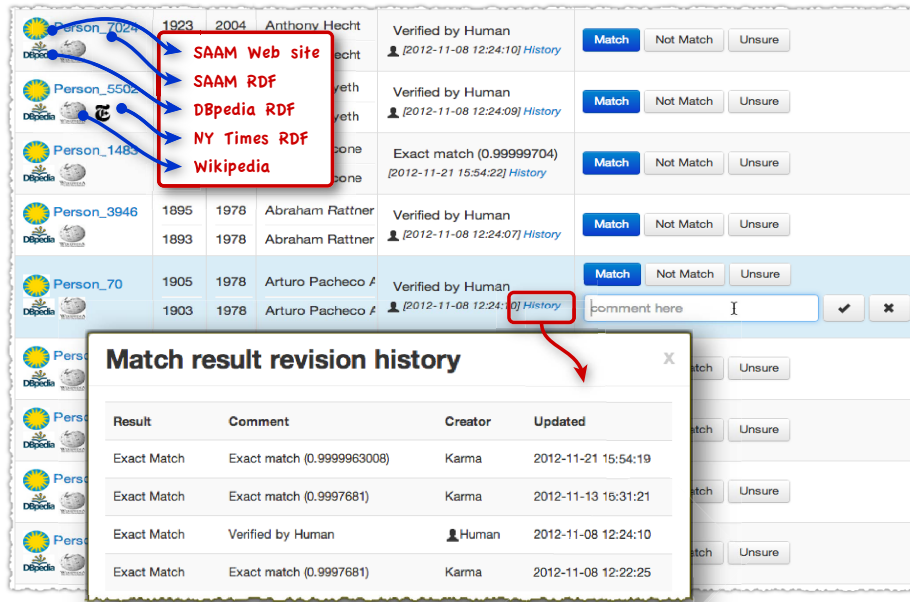


Fig. 6. The KARMA interface enables users to review the results of linking.

Users review the links using the Web interface depicted in Figure 6. The interface is a visualization and editor of the underlying PROV RDF records. Each row represents a link. The first cell shows the records being linked: the top part shows links to information about the SAAM record and the bottom part shows links to information for a record in an external source. The next columns show the data values that were used to create the link and information about its revision history. The last column shows buttons to enable users to revise links and provide comments (recorded as PROV records). SAAM personnel used this interface to verify all links in our ground truth dataset.

Lesson 7: PROV is a suitable technology for curating the links. In addition to supporting the user interface for human verification of links, the PROV representation affords other benefits. We can use SPARQL statements to construct a dataset of owl:same-as triples containing only those links that have been verified by a human (suitable for publication on the Smithsonian Web site) or a dataset containing all links with a matching score above a given threshold (suitable for other applications). Similarly, when the underlying data changes (e.g., there is a new version of DBpedia) or a new version of the matching software becomes available, it is possible to retrieve the affected links.

5 Related Work

There has been a great deal of recent interest in publishing museum data as Linked Open Data. One of the most ambitious efforts is a European project,

called Europeana, which has published the metadata on 17 million items from 1,500 cultural institutions [7]. This project developed a very rich and comprehensive ontology, called the Europeana Data Model (EDM) and then standardized the data that each organization can publish, which is how they were able to aggregate data from such a large number of cultural institutions. The focus of that effort was on developing a comprehensive data model and mapping all of the data to that model. Several smaller efforts focused on the problem of how to map rich metadata into RDF and preserve the content of the original data. This includes the MuseumFinland, which published the metadata on 4,000 cultural artifacts from museums and sites in Finland [8] and the Amsterdam Museum [3], which published the metadata on 73,000 objects from that museum. In both of these efforts the data is first mapped directly from the raw source into RDF and then there are complex mapping rules to transform the RDF into an RDF expressed in terms of their chosen domain ontology. The actual mapping process requires the use of Prolog rules for some of the more complicated cases. Finally, there is another effort, called the LODAC Museum, which published the metadata from 114 museums and research institutes in Japan. In that effort, they defined a relatively simple ontology that consists of objects, artists, and institutions to simplify the mapping process.

In our work on mapping the 41,000 objects from SAAM, we went beyond the previous work in several important ways. First, in terms of the mapping process, we developed an approach that supports the mapping of complex sources (both relational and hierarchical) into a rich domain ontology [9]. This approach is in contrast to previous work, which first maps the data directly into RDF [2] and then aligns the RDF with the domain ontology [1]. As described earlier, we build on the EDM ontology developed as part of the Europeana effort, which is a fairly rich domain ontology that is easily extensible. Our approach makes it possible to preserve the richness of the original metadata sources, but unlike the MuseumFinland and the Amsterdam Museum projects, a user does not need to learn a complex rule language and does not need to do any programming.

Second, in terms of the linking process, we have performed significantly more linking of data than any of these previous efforts. There has been a significant body of literature on linking data across sources and the most closely related work on linking is the work on Silk [14] and the work on entity coreference in RDF graphs [13]. Silk provides a nice framework that allows a user to define a set of matching rules and weights that determine whether two entities should be matched. We tried to use Silk on this project, but we found it extremely difficult to write a set of matching rules that produced high quality matches. The difficulty was due to a combination of missing data and the variation in the discriminability of different data values. The approach that we used in the end was inspired by the work on entity conference by Song and Heflin [13], which deals well with missing values and takes into account the discriminability of the attribute values in making a determination of the likelihood of a match.

Third, because of the importance to the Smithsonian of producing a high-quality set of linked data, we developed a curation tool that allows an expert from

the museum to review and approve or reject the links produced automatically by our system. Previous work has largely ignored the issue of link quality (Halpin et al. [6] reported that in one evaluation roughly 51% of the same-as links were found to be correct). The exception to this is the effort by the NY Times to map all of their metadata to linked data through a process of manual curation. In order to support a careful evaluation of the links produced by our system, we developed the linking approach that allows a link reviewer to see the data that is the basis for the link and to be able to drill down into the individual sources to evaluate a link.

6 Conclusions and Future Work

In this paper we described our work on mapping the data of the Smithsonian American Art Museum to Linked Open Data. We presented the end-to-end process of mapping this data, which includes the selection of the domain ontologies, the mapping of the database tables into RDF, the linking of the data to other related sources, and the curation of the resulting data to ensure high-quality data. This initial work provided us with a much deeper understanding of the real-world challenges in creating high-quality link data.

For the Smithsonian, the linked data provides access to information that was not previously available. The Museum currently has 1,123 artist biographies that it makes available on its website; through the linked data, we identified 2,807 matches to DBpedia. They can now embed the Wikipedia biographies into their collection information, increasing the biographies they offer by 60%. Via the links to DBpedia, they now have links to the New York Times, which includes obituaries, exhibition and publication reviews, auction results, and more. They can embed this additional rich information into their records, including 1,759 Getty ULAN identifiers, to benefit their scholarly and public constituents.

The larger goal of this project is not just to map the SAAM data to Linked Open Data, but rather to develop the tools that will enable any museum or other organization to map their data to Linked Data themselves. We have already developed the KARMA integration tool, which greatly simplifies the problem of mapping structured data into RDF, a high-accuracy approach to linking datasets, and a new curation tool that allows an expert to review the links across data sources. Beyond these techniques and tools, there is much more work to be done. First, we plan to continue to refine and extend the ontologies to support a wide range of museum-related data. Second, we plan to continue to develop and refine the capabilities for data preparation and source modeling in KARMA to support the rapid conversion of raw source data into RDF. Third, we plan to generalize our initial work on linking data and integrate a general linking capability into KARMA that allows a user to create high-accuracy linking rules and to do so by example rather than having to write the rules by hand.

We also plan to explore new ways to use the linked data to create compelling applications for museums. A tool for finding relationships, like EverythingIsConnected.be [12], has great potential. We can imagine a relationship finder appli-

cation that allows a museum to develop curated experiences, linking artworks and other concepts to present a guided story. The Museum could offer pre-built curated experiences or the application could be used by students, teachers, and others to create their own self-curated experiences.

References

1. Bizer, C., Schultz, A.: The R2R Framework: Publishing and Discovering Mappings on the Web. In: Proceedings of the First International Workshop on Consuming Linked Data (2010)
2. Bizer, C., Cyganiak, R.: D2R Server—publishing relational databases on the semantic web. In: Poster at the 5th International Semantic Web Conference (2006)
3. de Boer, V., Wielemaker, J., van Gent, J., Hildebrand, M., Isaac, A., van Ossenburg, J., Schreiber, G.: Supporting linked data production for cultural heritage institutes: the amsterdam museum case study. *The Semantic Web: Research and Applications* pp. 733–747 (2012)
4. Cohen, W., Ravikumar, P., Fienberg, S., et al.: A comparison of string distance metrics for name-matching tasks. In: Proceedings of the IJCAI-2003 Workshop on Information Integration on the Web (IIWeb-03). pp. 73–78 (2003)
5. Goel, A., Knoblock, C.A., Lerman, K.: Exploiting Structure within Data for Accurate Labeling Using Conditional Random Fields. In: Proceedings of the 14th International Conference on Artificial Intelligence (ICAI) (2012)
6. Halpin, H., Hayes, P., McCusker, J., McGuinness, D., Thompson, H.: When owl: sameas isnt the same: An analysis of identity in linked data. Proceedings of the 9th International Semantic Web Conference pp. 305–320 (2010)
7. Haslhofer, B., Isaac, A.: data.europeana.eu: The europeana linked open data pilot. *International Conference on Dublin Core and Metadata Applications* (2011)
8. Hyvönen, E., Mäkelä, E., Salminen, M., Valo, A., Viljanen, K., Saarela, S., Junnila, M., Kettula, S.: Museumfinland: Finnish museums on the semantic web. *Web Semantics* 3(2), 224–241 (2005)
9. Knoblock, C., Szekely, P., Ambite, J.L., Goel, A., Gupta, S., Lerman, K., Muslea, M., Taheriyan, M., Mallick, P.: Semi-Automatically Mapping Structured Sources into the Semantic Web. In: Proceedings of the 9th Extended Semantic Web Conference (ESWC) (2012)
10. Lafferty, J., McCallum, A., Pereira, F.: Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In: Proceedings of the 18th International Conference on Machine Learning (2001)
11. Matsumura, F., Kobayashi, I., Kato, F., Kamura, T., Ohmukai, I., Takeda, H.: Producing and consuming linked open data on art with a local community. In: Proceedings of the Third International Workshop on Consuming Linked Data (COLD2012). vol. 905. *CEUR Workshop Proceedings* (2012)
12. Sande, M.V., Verborgh, R., Coppens, S., Nies, T.D., Debevere, P., Vocht, L.D., Potter, P.D., Deursen, D.V., Mannens, E., , de Walle, R.V.: Everything is connected. In: Proceedings of the 11th International Semantic Web Conference (ISWC) (2012)
13. Song, D., Heflin, J.: Domain-independent entity coreference for linking ontology instances. *ACM Journal of Data and Information Quality (ACM JDIQ)* (2012)
14. Volz, J., Bizer, C., Gaedke, M., Kobilarov, G.: Silk—a link discovery framework for the Web of Data. In: Proceedings of the 2nd Linked Data on the Web Workshop. pp. 559–572 (2009)