

Exploring Linked Data For The Automatic Enrichment of Historical Archives

Gary Munnely, Harshvardhan J. Pandit, and Séamus Lawless

Adapt Centre, Trinity College Dublin, Dublin, Ireland
{firstname.lastname}@adaptcentre.ie

Abstract. With the increasing scale of online cultural heritage collections, the efforts of manually adding annotations to their contents become a challenging and costly endeavour. Entity Linking is a process used to automatically apply such annotations to a text based collection, where the quality and coverage of the linking process is highly dependent on the knowledge base that informs it. In this paper, we present our ongoing efforts to annotate a corpus of 17th century Irish witness statements using Entity Linking methods that utilise Semantic Web techniques. We discuss problems faced in this process and attempts to remedy them.

Keywords: entity linking, ontology creation, automatic enrichment

1 Introduction

The promise of Semantic Web [1] is an attractive one for any individual who is involved in cultural heritage research. It is a promise of powerful search, seamless integration and informed, reasoned decision making between the many siloed instances of data which prevail across domains. Unfortunately, before a collection can take advantage of the benefits gained by being a part of the Semantic Web, it must be annotated with a suitable vocabulary. This annotation process can be an expensive and challenging task for a number of reasons, not least of which is the time and labour cost of employing people to read and manually annotate the contents of the collection. Given the exponential effort of a manual annotation process, it is not surprising that there has been some interest in the effectiveness of applying automatic annotation tools to cultural heritage collections [2–4].

Of the variety of tools and methods that exist for extracting information from a collection, this paper focuses on those concerned with the problem of Entity Linking (EL) [5]. While EL has seen much research in recent years, the lack of suitable semantic web resources to inform the EL process is often a notable weakness which undermines efforts to apply EL methods to cultural heritage resources. To some extent we may accept this as an inevitable limitation due to the immense variety of cultural heritage collections that exist. It is improbable that we will ever have a single centralised source of information which covers all aspects of cultural heritage. However, as the scale of digitised cultural heritage collections grow (Europeana¹ alone currently curates more than 50 million

¹ <https://www.europeana.eu/portal/en>

different items and DPLA² hosts almost 21 million resources), it becomes increasingly worthwhile to consider how we might deal with these limitations for collections both great and small.

In this paper we present a discussion on the role of semantic web resources in the task of automatically enriching digitised cultural heritage collections using EL methods. Our discussion is motivated by ongoing efforts to annotate a collection of 17th century Irish witness statements so that they may be integrated as semantic web resources and avail of benefits such enrichment provides. We present some of the challenges faced and lessons learned in the course of this endeavour. We also describe a process by which we are attempting to construct a knowledge base for Entity Linking using record resolution methods. The contribution of this paper is to demonstrate and emphasise the importance of structure in Semantic Web resources. With due consideration it is possible to create new ontologies which may help to facilitate the EL process.

2 Related Work

2.1 Entity Linking

Entity Linking (EL) refers to a specific challenge in computer science whereby a series of unknown textual mentions of entities (commonly termed “surface forms”) are provided as input to a disambiguation service. The service is tasked with mapping each of the surface forms to an unambiguous referent entity. To provide a concrete example, given the input sentence, “*I Henry Jones Doctor in Divinity in obedience to his majesties Commission...*” and a request to identify the entity “Henry Jones”, an EL service might return a reference to the URI [http://dbpedia.org/page/Henry_Jones_\(bishop\)](http://dbpedia.org/page/Henry_Jones_(bishop)), identifying the subject of the reference as the 17th century Anglican Bishop, as opposed to the fictional character played by Séan Connery in the 1989 film “Indiana Jones and the Last Crusade”.

In order to perform this mapping process, an EL system fundamentally requires two components:

1. a knowledge base that stores information about all the entities of which the system is aware, and
2. a referent selection method, which uses evidence extracted from the knowledge base and present in any prevailing information surrounding the surface form to arrive at a set of likely referents for each ambiguous mention.

Given an ambiguous set of mentions, the EL system retrieves from the knowledge base, a set of candidate referents to which an entity mention may be referring. This is usually based on some fuzzy retrieval method. A variety of heuristics are applied and the system eliminates candidate referents which are unlikely to be the subjects of the mentions. Eventually it arrives at a set of mappings from

² <https://dp.la/>

textual mentions to knowledge base URIs which unambiguously identifies the referents.

Numerous different methods and approaches to EL may be freely found in the literature [6–8]. Almost universally, these methods use some form of graph based measure as one of the heuristics in the referent selection process. After the candidates have been retrieved from the knowledge base, a graph derived from the relationships between the entities may be constructed. The nature of these relationships varies, but usually it is based on links between corresponding Wikipedia pages. If a strong network exists between a number of candidate referents, then there is a good chance that they are the correct disambiguation choice for the given set of mentions.

It is also common to augment the graph weights using contextual cues derived from the words surrounding an entity mention [8,9]. We note this because systems which consider this feature are based on the assumption that some contextual description of each entity exists in the knowledge base.

The structure and content of the knowledge base is crucial, not only for informing the disambiguation service that an entity exists, but also for providing information which helps the the disambiguation algorithm to distinguish good referents from poor referents. Many modern systems make use of DBpedia³ [10] and YAGO⁴ [11] for this task. These are a good choice for most problems due to the prevalence of links between entities and the long form descriptions of entities obtained from their corresponding Wikipedia articles. However, for cultural heritage collections it is often the case that the range of information contained in these Semantic Web resources are not complete enough to capture the variety of entities we see in cultural heritage collections.

EL systems have the potential to be extremely helpful when enriching cultural heritage collections with semantic data. These fully automated systems are capable of deducing suitable annotations for raw, flat, textual documents based on information that is fed to them via a knowledge base. It is easy to see how a suitably informed EL system might dramatically ease the process of semantically linking new cultural heritage artifacts as they are digitised.

Further discussion could be had surrounding the precise point in the digitisation process at which EL is applied. Are we linking metadata which has already been normalised by an expert, or is the system capable of dealing with the noisy, original, primary source content from which the digital artifact is derived? In the case of the latter, how does the system manage archaic references, evolving entities and other such anomalies present in the source collection?

2.2 Automatic Enrichment in Cultural Heritage

There have been a number of efforts to investigate the effectiveness of EL methods in the automatic enrichment of cultural heritage collections.

³ <http://wiki.dbpedia.org/>

⁴ <https://www.mpi-inf.mpg.de/departments/databases-and-information-systems/research/yago-naga/yago/#c10444>

A Europeana led task force produced a series of reports in 2015 which document their experience with evaluating different cultural heritage enrichment services and sourcing different descriptive vocabularies as targets for the annotation process. Their focus was on annotating metadata for digitised artifacts. The content of this metadata ranged from specific fields comprised of a single entity e.g. `dc:creator`, `dc:publisher` to more general, free-form data such as `dc:description`.

As part of the investigation, a comparative evaluation of seven cultural heritage EL services was conducted [12]. Each service used a different vocabulary for enrichment, however the investigators were able to normalise the annotations by exploiting the fact that many of them made reference to corresponding DBpedia and Geonames entities. Using an evaluation dataset that was developed based on a combination of automatic enrichment tools and manual human investigation the report showed that the accuracy of the targeted EL tools was extremely high for the chosen collections.

However, as a variety of previous studies have shown, while the accuracy of EL methods may be high, quite often only a very small percentage of entities contained in cultural heritage datasets may actually be linked with a referent. A recent study we performed on the 1641 depositions (see Section 3) showed that a human annotator could only identify referents for 33% of the people and locations in the depositions [13]. We would compare this to efforts by other scholars such as Agirre [14], who attempted to link Europeana artifacts to Wikipedia articles and discovered that only 22% of entities that he identified could be annotated in this manner. This is an important limitation of which we must be aware.

One aspect of the problem is simply that cultural heritage collections are so incredibly diverse, complicated and unique that finding a suitable Semantic Web resource with adequate coverage for all purposes is nigh impossible. This presents the question, how should we annotate a cultural heritage collection when an appropriate Semantic Web resource cannot be found? Moreover (and of particular importance to our own research) how should these new Semantic Web resources be structured in order to aid the automatic enrichment process?

Of particular note for this discussion is the work of Brando et al. [15] on the REDEN project which investigated methods of using multiple knowledge bases for disambiguation. This is an interesting approach which may help to fill the gaps in popular knowledge bases using the information contained in more tailored ones. In their experiments DBpedia was used in conjunction with the Bibliothèque Nationale de France (BnF) ontology on a collection of French literary works.

REDEN's candidate selection phase is based on a literal string comparison between the surface form and entities in the knowledge base. All candidates from all source ontologies are retrieved and a resolution step based on `owl:sameAs` and `skos:exactMatch` properties resolves duplicate mentions into a single reference. Once the candidates have been appropriately pruned, a degree centrality measure is used to select the referents.

REDEN demonstrated that developing EL methods which can avail of multiple knowledge bases may help with poor coverage, but this requires that it be possible to establish reliable, accurate mappings between ontologies. Indeed, this property also facilitated the evaluation conducted by the Europeana task force. This is an important consideration when developing new vocabularies for cultural heritage collections.

3 The 1641 Depositions

Our own research has focused on attempts to automatically annotate a collection of 17th century manuscripts using EL methods. This has been extremely challenging for a variety of reasons. However, we believe our experiences are a reasonably typical example of problems faced in this field.

The 1641 depositions are a collection of letters and witness statements taken from the people of Ireland during the 1641 Irish rebellion. The physical manuscripts are comprised of approximately 19,000 pages bound in 31 volumes. Ireland in 1641 was a tumultuous place, and while the accuracy of some of the witness statements may be questionable, the depositions provide an unparalleled window into this dark chapter in Irish history.

The depositions have been digitised, transcribed and annotated by a team of historical scholars who extracted references to people and locations, tagged depositions based on the nature of their contents, and preserved as much information about the physical manuscripts as possible including margin notes, original spelling etc. The resulting documents are stored in a combination of TEI annotated files and an SQL database. This data rich digital resource presents many interesting and exciting opportunities for computer scientists to begin experimenting with methods of analysing and extracting new information from this historical collection.

Working with the digital versions of the depositions comes with a number of challenges, not least of which is the inconsistent nature of the spelling and grammar used throughout. English was still a developing language in 1641, which means that a vast array of variant spellings for names and common words exist across the documents. The below extract from *The Deposition of Phillip Sergeant*⁵ provides an example of these anomalies:

“And by those faire promisses the said ffitzpatrick getting possession both of their persons & goodes, they there behoulding daily cruelties & murders vpon other English and belike suspecting the like to be exercised against themselues, desired fled away secretly ~~o-n-to~~ to Mountrath”

From the historians’ work, we find that the depositions contain references to more than 60,000 people and 7,000 locations. The people in question range from individuals of great historical importance such as Sir Oliver Cromwell, Sir Phelim O’Neill, and King Charles I, to individual servants and common folk

⁵ <http://1641.tcd.ie/deposition.php?depID=815351r406>

who were affected by the rebellion. Locations similarly range from cities such as Dublin which still flourish today, to small plots of land which have been lost either as their names changed or borders shifted.

We know that several of the entities extracted from the depositions are duplicates. However, the huge range in spelling variations and naming conventions makes it extremely difficult to determine which mentions of entities in the depositions might be references to the same person or place. Compounding this problem is the fact that the severity of textual noise means that standard NLP tools can struggle with simpler tasks such as sentence chunking or Named Entity Recognition. Performing reliable analysis based on the language of the depositions is, to say the least, difficult.

We are not the first to attempt to decipher the contents of the depositions using computational methods. The CULTURA project [16] developed and applied a range of tools to provide a personalised experience for individuals who are interested in exploring the collection. This project was extremely successful and produced a number of valuable utilities for working with collections of this nature. However, the depositions' content remains in its original SQL database, disconnected from the Semantic Web.

An earlier study conducted on a manually annotated subset of the depositions attempted to assess the feasibility of automatically enriching the collection using standard Entity Linking tools [13]. From this study it was shown that only 33% of entities in the annotated subset had a corresponding referent in DBpedia. It was also observed that, of the eight Entity Linkers evaluated, no single tool could satisfactorily annotate the test corpus. Individually some did show promise on specific aspects of the linking problem, but these were undermined by weaknesses elsewhere. For example, AGDISTIS [7] often correctly abstained from annotating where no referent existed in DBpedia. Yet when considering the accuracy of the systems on entities that should have been annotated, the highest performing system achieved an F1 score of only 0.33, indicating that EL systems are greatly challenged by the depositions.

To some extent, the problems faced by the EL systems can be ascribed to the knowledge base. First, it is clear that the coverage of DBpedia is not sufficient due to the vast proportion of entities in the gold standard which were given a NIL label. It is also worth investigating whether a more tailored knowledge base might improve the F1 score of annotation systems on entities which should have been annotated.

4 Identifying Candidate Knowledge Bases

When investigating Entity Linking for the depositions, we focus on people and places due to their perceived importance in the documents. For historians there are still a number of unanswered research questions about the motivations and influences behind the rebellion. A linked data solution may assist them when investigating these questions. However, modelling the depositions is extremely challenging for a number of reasons.

First, there is no true, definitive list of people and places on which to base the ontology. Ideally any ontology that we create would be populated with a set of distinct entities that can be found in the depositions. While there are resources which can help us to determine this set of entities (as discussed below), there is still noise present in these sources. Sometimes people are referred to by lineage rather than their actual name, e.g. “the heirs of Mr. Gale”, or even by title e.g. “Bishop of Meath”. Given these ambiguities, there is much risk of accidentally omitting or conflating entities when the ontology is being constructed.

Second, the inconsistent language of the depositions means that multiple variant spellings for people and places can be found throughout the collection. If a suitable ontology can be constructed to represent each entity, discovering all the possible variant names by which it may be referenced would be a monumental task. Sometimes these variations are minor spelling differences e.g. “Florence FitzPatrick” being referred to as “Fflorenc Ffitz Patrick”, but some are more severe, such as the “Barony of Fassadinin” being referred to as the “Barrony of ffaassa and Dyninge”. Detecting such differences is difficult through an automated process and requires an expert to assess its correctness.

Third, if we are to construct this ontology with an eye to automatic enrichment, then the inclusion of links between entities and how to establish them is an important consideration. We could use familial connections, but we are not aware of any reliable sources which document these in a readily adoptable manner. On what basis then are we to establish relationships between our entities? Currently, there is no reliable way to specify that a relation is likely without stating it as a fact in an ontology.

We must also exercise some degree of caution in our attempts to annotate the depositions. If the intention is to assist scholars with their research, then the information conveyed by the proposed solution must be accurate. This can be a subtle problem. For example, if we consider the entity “the Pope”, should this be used to describe the role of the head of the Catholic Church, or should it describe an individual who held that role? If we assume the latter, then we must be sure to refer to the correct pope for the source document, which involves additional knowledge that may not be readily available in the knowledge base. Pope Urban VIII held the position until 1644 when he passed away and was replaced by Pope Innocent X. Modelling evolving entities such as these is a common problem in the cultural heritage domain.

In spite of these challenges, resources do exist which can help us to generate lists of distinct entities. Three primary sources at our disposal are:

- The Down Survey: A complete national survey of land in Ireland after the rebellion. The survey was conducted in order to establish which lands should be forfeited as penalty for crimes during the rebellion.
- The Statute Staple: A record of transactions between individuals. The staple documents goods bought and sold, and provides information about debts owed between various parties before the rebellion

- The Books of Survey and Distribution: A list of properties held by various land owners. These documents were used to determine taxes based on land ownership.

These documents have been the subject of historical research for a number of years and were some of the major contributing sources for the Petty Maps project⁶. Notably, The Down Survey and The Books of Survey and Distribution provide lists of important land owners resident in Ireland during the 1641 rebellion, possibly providing a definitive list of both people and locations. Statute Staple may be used to reveal relationships between these entities. We believe that it is possible to structure this information such that an EL system may use it for automatic enrichment of Irish cultural heritage resources.

We also note that two secondary sources which may be helpful are the Oxford Dictionary of National Biography (ODNB)⁷ and the Dictionary of Irish Biography (DIB)⁸. These resources are comprised of a number of biographies about significant figures in the history of the British Isles. These resources have the advantage that they are better structured than the primary sources identified above, but they are not as complete with respect to the entities that interest us. We have conducted a separate investigation into the construction of knowledge bases using these resources [17] as they present a different set of challenges to the three primary sources.

5 Resolving Entities Across Sources

Given the three primary sources described in Section 4, we have begun the process of constructing an ontology to model the entities present in the depositions. In order to facilitate EL methods, we are attempting to capture features that are commonly used by EL algorithms.

In the simplest terms, our objective is to construct a list of unique entities which are present in each of the records. Such a list would form the basis of a knowledge base by providing a set of entities which we expect to find in the depositions. Through further examination of the available sources e.g. Statute Staple, we may extract additional information about entities, such as relationships, providing more evidence which an EL system might use when investigating the content of a collection.

Effectively, this is a record resolution problem. We have a set of disparate records and our objective is to determine which records refer to the same individuals. The initial phase of this resolution was performed manually by a team of historians working with The Down Survey. The historians extracted two lists of unique landowners for the periods surrounding 1641 and 1670. They also extracted a list of townlands and parishes which were annotated with corresponding longitude and latitude coordinates. This provides a reliable foundation on which

⁶ <http://downsurvey.tcd.ie/index.html>

⁷ <http://www.oxforddnb.com>

⁸ <http://dib.cambridge.org/>

we can build our knowledge base, but the list is known to be incomplete. In particular, because the focus of the list of people is on landowners, many of the more common individuals in the depositions (servants, etc.) are not present.

Given the historians' lists of landowners, we have attempted to identify instances of these entities in Statute Staple and The Books of Survey and Distribution. There is often little evidence available for this process beyond the name of an individual. We have found that the Fellegi-Sunter method of comparing records [18] is an effective means of filtering candidate resolutions to a manageable pool. However, ultimately, a manual check by a trained historian is required to select which record resolutions are valid.

Gradually this process yields a list of unique entities, a variety of surface forms by which they may be referenced and relationships which exist between the entities in question. Note that we are not necessarily concerned with the nature of a relationship from the perspective of EL. Most Entity Linking algorithms which use relationships between entities are only concerned with the binary presence or absence of a connection. For the purposes of transparency, we state that the relationships in our knowledge base are derived from financial records documenting debts between parties, but this information is not captured by the knowledge base we have created.

We have made use of the DBpedia and FOAF vocabularies to model the properties of individuals extracted thus far. The information available to us is reasonably simple and these vocabularies are adequate for capturing properties such as names of an individual. The DBpedia ontology's `dbo:related` property is used to capture relationships between entities.

The nature of the depositions themselves yields information about the lifespan of entities in our knowledge base. While we do not know the precise dates of birth and death for the entities extracted from the records, we do know that they were alive in 1641. This is analogous to the concept of "floruit", which is essentially a fuzzy time period during which an individual is known to have existed. The vocabularies chosen so far are not adequate for capturing this uncertainty. However, we have found that CIDOC-CRM's timespans allow us to express uncertainty around an individual's lifespan. Although not a typically feature of EL systems, capturing temporal information such as this makes it possible to quickly filter a knowledge base to a range of feasible referents for a given surface form given the period of the text from which it was extracted. Hence it is of benefit to capture some aspect of the temporal properties of entities in the depositions, however crude the representation may be.

This process of constructing a knowledge base is admittedly slow. What we have found is that the features we must extract in order to inform an EL system are reasonably simple – surface forms and relationships. Context vectors for an entity are also a common feature for many EL systems, but given the linguistic anomalies of the depositions this is unlikely to be a helpful property to capture. It is frustrating to observe in the case of the depositions that the knowledge we require exists in digital format. It is simply distributed across a number of disparate repositories. Constructing semantic resources which structure this data

requires us to tackle record linkage problems before we can focus on the task of automatic enrichment. We are confident, however, that the resulting resource will be useful for EL, not only on the depositions, but on other cultural heritage archives for the British Isles.

6 Discussion

Performing automatic enrichment of cultural heritage collections is challenging for a variety of reasons. As evidenced by our own experience, and the documented experience of other researchers, finding knowledge bases with adequate coverage for a given cultural heritage resource is extremely difficult. While developing an entirely new ontology that does not reuse existing knowledge is a solution, if not done properly it can lead to inaccurate or incompatible knowledge representations that negate one of the greatest benefits of linked data i.e. connectivity among disparate collections. It is of far greater benefit to the community if these new vocabularies can be integrated with existing semantic web resources in a seamless fashion.

Due to the issue with well-known knowledge bases not covering a large percentage of the entities in specialised cultural heritage collections, it is likely that curators of such resources will need to develop their own ontologies in order to accurately represent the semantics of their data. While it is good to expand the web of knowledge with this new information, we suggest that due care be given to the structure of these resources and to how this structure may lend itself to informing automatic enrichment processes going forward. Methods such as REDEN may exploit `owl:sameAs` or similar relationships between a new ontology and more established ones in order to knit together various knowledge bases for the EL process. If automatic enrichment services can make use of the information in new linked data resources, then future annotation processes may be expedited as new collections are digitised and made available.

At present, the ontology we are constructing is disconnected from the greater semantic web. Our focus has been on the resolution of entities across the resources available to us. However, once this process is completed and important next step will be to associate entities in our knowledge base with their corresponding entities in more established knowledge bases such as Geonames or DBpedia.

Acknowledgements

The ADAPT Centre for Digital Content Technology is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

References

1. Berners-Lee, T., Hendler, J., Lassila, O.: The semantic web. *Scientific american* **284**(5) (2001) 34–43

2. Van Hooland, S., De Wilde, M., Verborgh, R., Steiner, T., Van de Walle, R.: Exploring entity recognition and disambiguation for cultural heritage collections. *Digital Scholarship in the Humanities* **30**(2) (2015) 262–279
3. Wilde, M.D.: Improving Retrieval of Historical Content with Entity Linking. In Morzy, T., Valduriez, P., Bellatreche, L., eds.: *New Trends in Databases and Information Systems. Communications in Computer and Information Science*, Springer International Publishing (September 2015) 498–504
4. Stiller, J., Petras, V., Gäde, M., Isaac, A.: Automatic enrichments with controlled vocabularies in europeana: Challenges and consequences. In: *Euro-Mediterranean Conference*, Springer (2014) 238–247
5. Shen, W., Wang, J., Han, J.: Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering* **27**(2) (2015) 443–460
6. Ganea, O.E., Ganea, M., Lucchi, A., Eickhoff, C., Hofmann, T.: Probabilistic bag-of-hyperlinks model for entity linking. In: *Proceedings of the 25th International Conference on World Wide Web, International World Wide Web Conferences Steering Committee* (2016) 927–938
7. Usbeck, R., Ngomo, A.C.N., Röder, M., Gerber, D., Coelho, S.A., Auer, S., Both, A.: Agdistis-graph-based disambiguation of named entities using linked data. In: *International Semantic Web Conference*, Springer (2014) 457–471
8. Yosef, M.A., Hoffart, J., Bordino, I., Spaniol, M., Weikum, G.: Aida: An online tool for accurate disambiguation of named entities in text and tables. *Proceedings of the VLDB Endowment* **4**(12) (2011) 1450–1453
9. Zwicklbauer, S., Seifert, C., Granitzer, M.: Robust and collective entity disambiguation through semantic embeddings. In: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '16, New York, NY, USA, ACM* (2016) 425–434
10. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morsey, M., Van Kleef, P., Auer, S., et al.: Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web* **6**(2) (2015) 167–195
11. Suchanek, F.M., Kasneci, G., Weikum, G.: Yago: A core of semantic knowledge. In: *Proceedings of the 16th International Conference on World Wide Web. WWW '07, New York, NY, USA, ACM* (2007) 697–706
12. Manguinhas, H., Freire, N., Isaac, A., Stiller, J., Charles, V., Soroa, A., Simon, R., Alexiev, V.: Exploring comparative evaluation of semantic enrichment tools for cultural heritage metadata. In: *International Conference on Theory and Practice of Digital Libraries*, Springer (2016) 266–278
13. Munnely, G., Lawless, S.: Investigating entity linking in early english legal documents. In: *Digital Libraries (JCDL), ACM/IEEE Joint Conference on.* (2018)
14. Agirre, E., Barrena, A., Lacalle, O.L.D., Soroa, A., Fern, S., Stevenson, M.: Matching Cultural Heritage items to Wikipedia. In: *LREC.* (2012) 1729–1735
15. Brando, C., Frontini, F., Ganascia, J.G.: REDEN: Named Entity Linking in Digital Literary Editions Using Linked Data Sets. *Complex Systems Informatics and Modeling Quarterly* (7) (July 2016) 60 – 80
16. Steiner, C.M., Agosti, M., Sweetnam, M.S., Hillemann, E.C., Orio, N., Ponchia, C., Hampson, C., Munnely, G., Nussbaumer, A., Albert, D., et al.: Evaluating a digital humanities research environment: the CULTURA approach. *International Journal on Digital Libraries* **15**(1) (2014) 53–70

17. Munnely, G., Lawless, S.: Constructing a knowledge base for entity linking on irish cultural heritage collections. In: Proceedings of the 14th International Conference on Semantic Systems. (in press)
18. DuVall, S.L., Kerber, R.A., Thomas, A.: Extending the fellegisunter probabilistic record linkage method for approximate field comparators. *Journal of Biomedical Informatics* **43**(1) (2010) 24 – 30