

A Web-based Geo-resolution Annotation and Evaluation Tool

Beatrice Alex, Kate Byrne, Claire Grover and Richard Tobin

School of Informatics

University of Edinburgh

{balex, kbyrne3, grover, richard}@inf.ed.ac.uk

Abstract

In this paper we present the Edinburgh Geo-annotator, a web-based annotation tool for the manual geo-resolution of location mentions in text using a gazetteer. The annotation tool has an inter-linked text and map interface which lets annotators pick correct candidates within the gazetteer more easily. The geo-annotator can be used to correct the output of a geoparser or to create gold standard geo-resolution data. We include accompanying scoring software for geo-resolution evaluation.

1 Introduction

Many kinds of digitised content have an important geospatial dimension. However not all geospatial information is immediately accessible, particularly in the case where it is implicit in place names in text. The process of geo-resolution (also often referred to as geo-referencing, geoparsing or geotagging) links instances of textual geographic information to location coordinates, enabling searching and linking of digital content using its geospatial properties.

Geo-resolution tools can never be completely accurate and their performance can vary significantly depending on the type and quality of the input texts as well as on the gazetteer resources they consult. For this reason, users of text collections are frequently disappointed in the results of geo-resolution and, depending on their application and dataset size, they may decide to take remedial action to improve the quality. The tool we describe here is a web-based, manual annotation tool which can be used to correct the output of geo-resolution. It has been developed in step with our geo-resolution system, the Edinburgh Geoparser (Grover et al., 2010), but it could also be used to correct the output of other tools. In our work, we use the geo-annotator to create gold-standard material for geo-resolution evaluation and have produced accompanying scoring software.¹

2 Related Work

Within the field of NLP, SpatialML is probably the best known work in the area of geo-referencing. SpatialML is an annotation scheme for marking up natural language references to places and grounding them to coordinates. The SpatialML corpus (Mani et al., 2008) instantiates this annotation scheme and can be used as an evaluation corpus for geo-resolution (Tobin et al., 2010). Other researchers develop their own geo-annotated corpora and evaluate against these, e.g. Clough (2005), Leidner (2007).

Within the field of Information Retrieval, there is an ACM special interest group on spatially-related information, SIGSPATIAL², with regular geographic IR conferences (GIR conferences) where geo-referencing research is presented, see for example Purves et al. (2007).

There are currently several geoparsing tools available, such as GeoLocate³, and CLAVIN⁴, as well as our own tool, the Edinburgh Geoparser. All of these enable users to geo-reference text collections but do

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Page numbers and proceedings footer are added by the organisers. Licence details: <http://creativecommons.org/licenses/by/4.0/>

¹The Edinburgh Geo-annotator will be available at <http://www.ltg.ed.ac.uk>.

²<http://www.sigspatial.org/>

³<http://www.museum.tulane.edu/geolocate/>

⁴<http://clavin.bericotechnologies.com/>

not address the question of how to interact with the geo-annotations in order to correct them, nor do they assist in creating evaluation materials for particular text collections.

The Edinburgh Geo-annotator has been developed in tandem with the Edinburgh Geoparser and earlier versions have been used in the GeoDigRef project (Grover et al., 2010) to create evaluation data for historical text collections as well as in the botanical domain (Llewellyn et al., 2012; Llewellyn et al., 2011) where we adapted it to allow curators to geo-reference the textual metadata associated with herbarium specimens. The current version has also been used to create gold standard data for Trading Consequences, a historical text mining project on mining location-centric trading information relevant to the nineteenth century (Klein et al., 2014). The Pelagios project, which deals with texts about the ancient world, has recently developed Recogito⁵, a geo-resolution correction tool similar to our own.

3 Annotation Tool

The Edinburgh Geo-annotator is a geo-resolution annotation tool which can be used to correct geo-resolution output or to create manually annotated gold standard data for evaluating geo-resolution algorithms and tools. The geo-annotator has a web-based interface allowing easy off-site annotation in inter-disciplinary projects by domain experts (who are not always necessarily the developers of the geo-referencing software).⁶ The interface allows users to select documents from a collection of prepared files containing annotated location entity mentions. By selecting and loading a document, the user can see its textual content and the location mentions highlighted within it.

The current tool is set up to select locations from a set of location candidates retrieved from GeoNames and visualised by pins on a Google Maps (v3) window. However, it can be configured to use candidates from a different location gazetteer. There are two files associated with each document: (1) an HTML file which contains the text of the document and (2) an XML file which contains the candidates for each location mention in the text and in which the annotations are stored. Candidates are linked to location mentions via identifiers.

All location mentions displayed in the text interface are highlighted in colour (see Figure 1). Those in red (e.g. **Dublin**) have one or more potential candidates in the gazetteer, while those in blue (e.g. **British Empire**) do not have candidate entries in the gazetteer. There are a number of reasons why a mention does not have a gazetteer entry. For example, the mention might be an old name of a location which is not stored in the gazetteer, or the mention contains a misspelling. During the annotation phase, the user is instructed to go through the red location mentions in the text and select the appropriate candidate.

In some cases there is only one candidate that can be selected (see Figure 2). The user can zoom to the correct location pin which when selected shows a popup with the relevant gazetteer information for that entry. The user can choose this candidate by pressing either “Select for this mention” if the choice is specific to the selected mention or “Select for all mentions” if the selection can be propagated for all mentions with the same string in the document. Once a pin is selected, it and the location mention in the text turn green. To undo a selection, the user can click on a green pin and press either “Deselect for this mention” or “Deselect for all mentions”.

In other cases, there are many candidates to choose from. For example, when clicking on the first location mention (Dublin) shown in Figure 1, the map adjusts to the central point of all 42 candidate locations. When reading a piece of text, human beings can often easily understand which location a place name refers to based on the context it appears in, which means that choosing between multiple candidates manually is not expected to be a difficult task. However, the number of location candidates that are suggested by GeoNames and consequently displayed in the interface can be limited in the data files, if for example the user only wants to choose between a small number of candidates.

In the case of Dublin (see Figure 1), the user would then zoom into the correct Dublin to select a candidate and discover that there are two pins which are relevant, *Dublin* – the capital, and *Baile Átha Cliath* – the Gaelic name for Dublin and its gazetteer entry referring to the administrative division (see Figure 3). The gazetteer information in the popup can assist the user to make a choice. In this case, it is clear from the context that the text refers to the capital. It might not always be as clearcut to choose

⁵<http://pelagios-project.blogspot.co.at/2014/01/from-bordeaux-to-jerusalem-and-back.html>

⁶The geo-annotator is run via a javascript programme which calls an update.cgi script on the server side to write the saved data to file. We have tested it in Safari, Firefox and Chrome.




Figure 1: When an example location mention (e.g. Dublin) is clicked the map adjusts to show all potential location candidates that exist in the gazetteer for this place name.

between multiple candidates. In such cases, it is important that the annotation guidelines provide detailed instruction as to which type of gazetteer entry to prefer.

If none of the candidates displayed on the map are correct, then the user must mark this by pressing “This mention” (or “All mentions”) in the box located at the top of right corner of the map (see Figure 1). Once there are only green or blue location mentions left in the text, the annotation for the selected document is complete and the user should press “Save Current Document” and move to the next document in the collection.

4 Geo-resolution Evaluation

It is important to be able to report the quality of a geo-resolver’s performance in concrete and quantifiable terms. Along with the annotation tool, we are therefore also releasing an evaluation script which compares the manually geo-resolved locations to those predicted by an automatic geoparser.⁷ We follow standard practice in comparing system output to hand-annotated gold standard evaluation data. The script evaluates the performance of the geo-resolution independently from geo-tagging, meaning that it only considers named entities which were tagged in the input to the manual geo-resolution annotation but not those that were missed. It is therefore preferable to use input data which contains manually annotated or corrected location mentions.

The evaluation script computes the number of correctly geo-resolved locations and accuracy in percent. Both figures are presented for a strict evaluation of exact match against gazetteer identifier and for a lax evaluation where the grid references of the gold and the system choice have to occur within a small distance of one another to count as a match. For a pair of location candidates (gold vs. system), we compute the Great-circle distance using a special case of the Vincenty formula which is most accurate for all distances.⁸ The lax evaluation is provided as even with clear annotation guidelines, annotators

⁷We provide Mac and Linux binaries of the evaluation scripts.

⁸For the exact formula, see: http://en.wikipedia.org/wiki/Great-circle_distance



Figure 2: Example candidate for the location mention *River Liffey* and its gazetteer entry information shown in a popup.




Figure 3: Choosing between multiple candidates for the same location mention.

can find it difficult to choose between different location types for essentially the same place (e.g. see the example for Dublin in Figure 3).

During the manual annotation, three special cases can arise. Some location mentions do not have a candidate in the gazetteer (those appearing in blue), while others do have candidates in the gazetteer but the annotator does not consider any of them correct. Occasionally there are location mentions with one or more candidates in the gazetteer but an annotator neither chooses one of them nor selects “none”. The latter cases are considered to be annotation errors, usually because the annotator has forgotten to resolve them. The evaluation excludes all three cases when computing accuracy scores but notes them in the evaluation report in order to facilitate error analysis (see sample output in Figure 4).

```
total: 11 exact: 10 (90.9%) within 6.0km 11 (100.0%)
note: no gold choice for British Empire
note: annotator selected "none" for Irish Free State
```

Figure 4: Sample output of the geo-resolution evaluation script. When setting the lax evaluation to 6km, one candidate selected by the system was close enough to the gold candidate to count as a match.

5 Summary

We have presented a web-based manual geo-resolution annotation and evaluation tool which we are releasing to the research community to facilitate correction of automatic geo-resolution output and evaluation of geo-resolution algorithms and techniques. In this paper we introduce the annotation tool and its main functionalities and describe two geo-resolution evaluation metrics with an example, namely strict and lax accuracy scoring. The release will contain more detailed documentation of the configuration and installation process and the document formats for the textual input and candidate gazetteer entries.

References

- Paul Clough. 2005. Extracting metadata for spatially-aware information retrieval on the internet. In *Proceedings of Workshop on Geographic Information Retrieval (GIR'05)*.
- Claire Grover, Richard Tobin, Kate Byrne, Matthew Woollard, James Reid, Stuart Dunn, and Julian Ball. 2010. Use of the Edinburgh geoparser for georeferencing digitised historical collections. *Phil. Trans. R. Soc. A*.
- Ewan Klein, Beatrice Alex, Claire Grover, Richard Tobin, Colin Coates, Jim Clifford, Aaron Quigley, Uta Hinrichs, James Reid, Nicola Osborne, and Ian Fieldhouse. 2014. Digging Into Data White Paper: Trading Consequences.
- Jochen L. Leidner. 2007. *Toponym Resolution in Text: Annotation, Evaluation and Applications of Spatial Grounding of Place Names*. Ph.D. thesis, School of Informatics, University of Edinburgh.
- Clare Llewellyn, Elspeth Haston, and Claire Grover. 2011. Georeferencing botanical data using text analysis tools. In *Proceedings of the Biodiversity Information Standards Annual Conference (TDWG 2011)*.
- Clare Llewellyn, Claire Grover, Jon Oberlander, and Elspeth Haston. 2012. Enhancing the curation of botanical data using text analysis tools. In Panayiotis Zaphiris, George Buchanan, Edie Rasmussen, and Fernando Loizides, editors, *Theory and Practice of Digital Libraries*, volume 7489 of *Lecture Notes in Computer Science*, pages 480–485. Springer Berlin Heidelberg.
- Inderjeet Mani, Janet Hitzeman, Justin Richer, Dave Harris, Rob Quimby, and Ben Wellner. 2008. SpatialML: Annotation scheme, corpora, and tools. In *Proceedings of the Sixth International Language Resources and Evaluation (LREC'08)*.
- Ross S. Purves, Paul Clough, Christopher B. Jones, Avi Arampatzis, Benedicte Bucher, David Finch, Gaihua Fu, Hideo Joho, Awase Khirmi Syed, Subodh Vaid, and Bisheng Yang. 2007. The design and implementation of SPIRIT: a spatially-aware search engine for information retrieval on the internet. *International Journal of Geographic Information Systems (IJGIS)*, 21(7).
- Richard Tobin, Claire Grover, Kate Byrne, James Reid, and Jo Walsh. 2010. Evaluation of georeferencing. In *Proceedings of Workshop on Geographic Information Retrieval (GIR'10)*.