

Automatic Construction of a Large-Scale Corpus for Geoparsing Using Wikipedia Hyperlinks

Keyaki Ohno¹ Hirotaka Kameko² Keisuke Shirai¹
Taichi Nishimura^{*3} Shinsuke Mori²

¹Graduate School of Informatics, Kyoto University, Japan

²Academic Center for Computing and Media Studies, Kyoto University, Japan

³LY Corporation, Japan

{ohno.keyaki.57r, shirai.keisuke.l82}@kyoto-u.jp

{kameko,forest}@i.kyoto-u.ac.jp tainishi@lycorp.co.jp

Abstract

Geoparsing is the task of estimating the latitude and longitude (coordinates) of location expressions in texts. Geoparsing must deal with the ambiguity of the expressions that indicate multiple locations with the same notation. For evaluating geoparsing systems, several corpora have been proposed in previous work. However, these corpora are small-scale and suffer from the coverage of location expressions on general domains. In this paper, we propose *Wikipedia Hyperlink-based Location Linking* (WHLL), a novel method to construct a large-scale corpus for geoparsing from Wikipedia articles. WHLL leverages hyperlinks in Wikipedia to annotate multiple location expressions with coordinates. With this method, we constructed the WHLL corpus, a new large-scale corpus for geoparsing. The WHLL corpus consists of 1.3M articles, each containing about 7.8 unique location expressions. 45.6% of location expressions are ambiguous and refer to more than one location with the same notation. In each article, location expressions of the article title and those hyperlinks to other articles are assigned with coordinates. By utilizing hyperlinks, we can accurately assign location expressions with coordinates even with ambiguous location expressions in the texts. Experimental results show that there remains room for improvement by disambiguating location expressions.

Keywords: Geoparsing, Corpus construction, Wikipedia

1. Introduction

Recognizing the spatial information indicated by location expressions in texts is a promising direction in text comprehension by machines. This enables us to automatically analyze large amounts of texts worldwide and read the texts about places of interest. In addition, it would be possible to organize texts based on their geographic characteristics.

Geoparsing is the task of estimating the latitude and longitude (coordinates) of location expressions in texts. Typically, geoparsing consists of two tasks: recognizing location expressions in texts (geotagging) and estimating their coordinates (geocoding). One of the challenges in geoparsing is the disambiguation of the expressions. For example, the word *Melbourne* is ambiguous as it could be a city in Australia, the United Kingdom, Canada, or the United States. We focus on such ambiguous location expressions where the same notation indicates multiple locations.

Previous work proposed corpora to evaluate the performance of disambiguation by geoparsing systems (Lieberman et al., 2010; Gritta et al., 2018b,a). These corpora are relatively small-scale (e.g., hundreds or thousands) or focus on specific domains. Thus, using these corpora suffers from the coverage of location expressions when evaluating the systems on general domains. Therefore, a large-

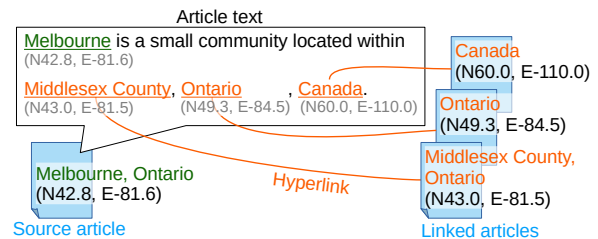


Figure 1: Coordinates annotation using hyperlinks. Location expressions with hyperlinks are annotated with the coordinates of the linked articles, and those equal to the article title are annotated with the coordinates of the article. A sample article of our corpus is shown in Figure 2.

scale corpus containing general articles is required to alleviate these issues.

In this paper, we propose *Wikipedia Hyperlink-based Location Linking* (WHLL) that automatically assigns multiple location expressions with coordinates in Wikipedia articles. Contrary to previous work (Gritta et al., 2018b), WHLL realizes the coordinates annotation to location expressions other than the article title by leveraging hyperlinks in Wikipedia dump files. Note that WHLL does not predict any coordinates and uses those registered in Wikipedia articles. Figure 1 shows a schematic of WHLL. Based on this method, we constructed the WHLL corpus, a new large-scale corpus for geop-

*Work done while the author was at Kyoto University.

arsing. It has more than 1.3M articles with about 7.8 unique location expressions per article, and 45.6% of total location expressions are ambiguous. Our code and data are publicly available.¹

2. Related Work

2.1. Methods for geoparsing

Various methods have been proposed in terms of disambiguation. This includes rule-based methods that refer to population, geographic distance between candidates or administrative level on a database (Aldana-Bobadilla et al., 2020), statistical methods that calculate probability distributions from data in advance (Speriosu and Baldrige, 2013; DeLozier et al., 2015), and machine learning-based methods (Awamura et al., 2015; Gritta et al., 2018a; Kulkarni et al., 2021).

2.2. Corpora for geoparsing

The evaluation of geoparsing, especially the task of estimating coordinates for each location expression in a text, has often relied on corpora such as Local-Global Lexicon (LGL) (Lieberman et al., 2010), Wikipedia Toponym Retrieval (WikToR) (Gritta et al., 2018b) and GeoVirus (Gritta et al., 2018a). **LGL** consists of 588 news articles. Each location expression was manually annotated. They collected newspapers located near locations whose names were ambiguous, such as the Paris News (Texas) and the Paris Post-Intelligencer (Tennessee). **WikToR** is a corpus created from 5,000 Wikipedia pages. They selected ambiguous location expressions using GeoNames and collected the Wikipedia first paragraphs of these locations using GeoNames API. Since each article in the corpus is annotated only with the coordinates of the article title, it has only one unique coordinate. **GeoVirus** consists of 229 WikiNews about global disease outbreaks and epidemics. Coordinates were manually determined, referring to Wikipedia or GeoNames.

In addition to evaluation data, machine learning-based geoparsing requires a considerable amount of training data. Some previous studies actually used more than 1 million Wikipedia pages for machine learning.

It is desirable that a corpus for geoparsing (i) contains many ambiguous location expressions and (ii) is unbiased in its sampling of locations. Furthermore, especially as training data, it is desirable that (iii) multiple location expressions appear within a text. This is because it is useful for learning the geographical relationships between each location expression. However, a large-scale corpus that

Corpus	#articles	#unique location expressions par article
LGL	588	4.2
WikToR	5,000	1.0
GeoVirus	229	6.6
WHLL corpus	1,315,117	7.8

Table 1: Comparison of the WHLL corpus and previously proposed corpora for geoparsing.

meets these requirements has not been published to date. Corpora such as those mentioned above are biased in their sampling, such as enriching ambiguous location expressions, for efficient evaluation. Therefore, we created a corpus that satisfies these three requirements. Table 1 shows that our corpus is large and contains ambiguous location expressions comparable to those of previous corpora.

2.3. Automatic dataset creation from Wikipedia metadata

Outside the field of geoparsing, there are various works on automatic dataset creation using Wikipedia metadata. Yang and Vozila (2014) utilized entity tags in Wikipedia as natural annotations for word segmentation tasks. Ling and Weld (2012) leveraged Wikipedia text and its hyperlinks to create training data for NER. Pan et al. (2017) proposed to link non-English name mentions to an English knowledge base automatically based on Wikipedia anchor links.

3. Wikipedia Hyperlink-based Location Linking

Wikipedia Hyperlink-based Location Linking (WHLL) is a novel method to automatically assign multiple location expressions in Wikipedia articles with coordinates. Here, we focus on articles of locations. Based on WHLL, we create a new corpus, the WHLL corpus. In this section, we first describe WHLL in Section 3.1 and then introduce the WHLL corpus and its statistics in Section 3.2.

3.1. Construction process

Many articles on Wikipedia are related to locations, and each article contains the single coordinates of the location. The articles contain more location expressions than the title, and several expressions are annotated with hyperlinks to other articles. Thus, location expressions in the article can be divided into two types: expressions with and without hyperlinks. WHLL automatically assigns these location

¹<http://www.lsta.media.kyoto-u.ac.jp/resource/data/WHLL/home-e.html>

Title	Melbourne, Ontario		
Text	<p>Melbourne is a small community located within Middlesex County, Ontario, Canada. It lies on the boundary between two municipalities, Strathroy-Caradoc and Southwest Middlesex. About half the population of Melbourne lives in each municipality. The community was probably named for Melbourne, Victoria, Australia.</p>		
Annotation	Span	Notation in the text	(Latitude, Longitude)
	(0, 9)	"Melbourne"	(42.81667, -81.55194)
	(46, 62)	"Middlesex County"	(43.00000, -81.50000)
	(64, 71)	"Ontario"	(49.25000, -84.50000)
	(73, 79)	"Canada"	(60.00000, -110.00000)
	(133, 150)	"Strathroy-Caradoc"	(42.95750, -81.61667)
	(155, 174)	"Southwest Middlesex"	(42.75000, -81.70000)
	(205, 214)	"Melbourne"	(42.81667, -81.55194)
	(280, 289)	"Melbourne"	(-37.81417, 144.96306)
(291, 299)	"Victoria"	(-37.00000, 144.00000)	

Figure 2: A sample article from our corpus (including 7 [Canadian](#) and 2 [Australian](#) location expressions). Annotation includes string span, notation, latitude, and longitude. The last [Melbourne](#) in the HTML text of this article had a hyperlink to another Wikipedia article, resulting in different coordinates annotated. Note that location expressions without hyperlinks (e.g., [Australia](#) in the sample) are ignored in the annotation.

expressions with coordinates of the original article and those of the linked articles.

The process consists of two steps: identifying the location expressions' span in the article and assigning coordinates to the identified expressions. For the expressions with hyperlinks, WHLL identifies the span of the expressions based on special tags (e.g., in HTML dump files, such expressions are surrounded by `<a>` tag). The expressions are then assigned the coordinates of the articles to which the expression hyperlinks.

For the expressions without hyperlinks, WHLL targets expressions of the article title and identifies their spans. More concretely, WHLL searches for expressions matching any of the following:

- The article title (e.g., "[Melbourne, Ontario](#)"),
- String obtained by removing the comma and the right side from the title (e.g., "[Melbourne](#)" from "[Melbourne, Ontario](#)"),
- String obtained by removing the parenthesis and its content from the title (e.g., "[Waterloo](#)" from "[Waterloo \(Albertson, North Carolina\)](#)").

The matched expressions are then assigned with the coordinates annotated on the current article.

WHLL utilizes two types of Wikipedia dump files: an HTML dump for obtaining formatted articles and a CirrusSearch dump for realizing the assignments of the expressions with coordinates based on hyperlinks. Since WHLL is automatic, it provides consistent annotations and can be applied to large texts with minimum cost. Further, one can update or newly create the corpus based on dump files. Note that WHLL constructs the corpus from Wikipedia articles; it does not newly predict coordinates.

3.2. The WHLL corpus

The WHLL corpus is based on two dump files: the Wikipedia HTML dump file released on July 1, 2023² and the Wikipedia CirrusSearch dump file released on July 10, 2023.³ We use only articles that are annotated with coordinates. Figure 2 shows a sample article in the corpus.

Table 2 shows the statistics of the WHLL corpus. The WHLL corpus has 1.3M articles and over 14.7M location expressions (1.6M unique expressions). 45.6% of the expressions are ambiguous: there are multiple coordinates corresponding to the single notation. In addition, 9.9% of the expressions are recessive: they are ambiguous expressions and do not refer to the most frequent coordinates. Disambiguating these expressions is challenging because simply choosing the most frequent coordinates fails to find the correct coordinates. Each article contains 7.8 unique location expressions on average, indicating more dense annotations than WikToR (Gritta et al., 2018b). As illustrated in Figure 3, the expressions in the WHLL corpus refer to locations around the world.

²https://dumps.wikimedia.org/other/enterprise_html/runs/20230701/enwiki-NS0-20230701-ENTERPRISE-HTML.json.tar.gz. Accessed on July 17, 2023 and checked the expiration on October 20, 2023.

³<https://dumps.wikimedia.org/other/cirrussearch/20230710/enwiki-20230710-cirrussearch-content.json.gz>. Accessed on July 11, 2023 and checked the expiration on October 20, 2023.

Total	
#articles	1,315,117
#sentences	23,334,035
#location expressions	14,726,908
(Ambiguous)	45.6%
(Ambiguous & Recessive)	9.9%
#unique location expressions	1,571,291
(Ambiguous)	8.1%
Per article	
#sentences	17.7
#tokens	420.1
#location expressions	11.3
#unique location expressions	7.8

Table 2: Statistics of our corpus. Ambiguous means the location expressions indicate multiple coordinates with the same notation.

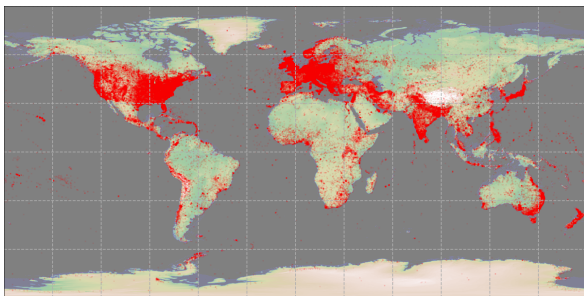


Figure 3: Coordinates that appear more than 10 times in the WHLL corpus.

4. Experiments

We conducted geocoding experiments to show the necessity of the disambiguation for the WHLL corpus. Here, we used ground-truth location expressions by assuming that geotagging was already finished.⁴ In this section, we first describe our method for geocoding and experimental settings and then report our results.

4.1. Method

We used GeoNames⁵ as a geographic database. Each entry in GeoNames has a name, a list of alternate names, coordinates and other information. We assigned an entry to each location expression by following strategies.

Familiarity-based strategy is the one that chooses an entry whose name matches the focused location expression and having the highest number of alternate names registered. This

⁴In previous work (Lieberman et al., 2010; Gritta et al., 2018a), geotagging was solved by using named entity recognition models.

⁵<https://www.geonames.org/>. Accessed on October 20, 2023.




Figure 4: Dependency-based strategy.

is based on our observation that famous locations tend to have many alternate names in GeoNames. If no candidates are found in the database, the coordinates just before in the sentence are copied.

Dependency-based strategy is the one that chooses an entry considering the dependencies between location expressions. We assume that the locations indicated by two dependent location expressions are geographically close. Using Universal Dependencies (UD) (McDonald et al., 2013) as the dependency structure, we selected coordinates by the following procedure (Figure 4):

- (1) Find mutually related location expressions using dependency structure between tokens.
- (2) For the modifier tokens, choose candidates (coordinates) in the same way as familiarity-based strategy. These location expressions tend to provide supplementary information in the sentence and indicate relatively large geographic locations.
- (3) For the head tokens, select the closest coordinates to (2) from the candidates or copying (2) if no candidates are found. These location expressions tend to provide the main information of the sentence and indicate relatively small geographical locations, often included in the location of (2).

4.2. Experimental settings

We used 10,000 randomly selected articles from the created corpus for evaluation. We used Stanford NLP Group’s stanza (Qi et al., 2020) as the tokenizer and the dependency parser (Dozat and Manning, 2017). Following previous work (Gritta et al., 2018a; Kulkarni et al., 2021), we used accuracy@161 km as an evaluation metric. This is the proportion of location expressions whose coordinates were estimated within 161 km of their actual coordinates.

4.3. Results

Figure 5 shows the results. The accuracy of the dependency-based strategy (0.58) is better than that of the familiarity-based one (0.56). This suggests that choosing the most famous locations does




Figure 5: Geocoding accuracy to tolerance error distance.

not always work well and that leveraging linguistic contexts by UD helps improve the accuracy of geocoding.

5. Conclusion

We have proposed Wikipedia Hyperlink-based Location Linking, a novel method to create a large-scale corpus for geoparsing leveraging hyperlinks. With this method, we have created a new corpus called the WHLL corpus. Each article in the WHLL corpus contains multiple location expressions annotated with coordinates. Our method and corpus enable researchers in this field to (i) train machine learning-based geoparsing models on a large-scale corpus and (ii) tackle geoparsing with ambiguous location expressions.

6. Limitation

In our corpus, not all the location expressions are necessarily annotated with coordinates. If a location expression in a Wikipedia page does not have a hyperlink, it is not annotated. Hyperlinks were often omitted for easy-to-understand expressions, such as country names listed alongside city names.

Also, regarding the location expression of the article title, coordinates annotation using 3.1 failed when the characters of the article title and those in the text were slightly different. In addition to variations in the writing of Wikipedia page authors, these differences were sometimes due to restrictions on the characters that can be used in article titles. For example, apostrophes in the article title had UTF-8 hexadecimal character code 27, but they sometimes appeared as CA BC in the text.

7. Acknowledgements

This work was supported by JSPS KAKENHI Grant Number JP 21H04376.

8. Bibliographical References

- Edwin Aldana-Bobadilla, Alejandro Molina-Villegas, Ivan Lopez-Arevalo, Shanel Reyes-Palacios, Victor Muñiz-Sanchez, and Jean Arreola-Trapala. 2020. [Adaptive geoparsing method for toponym recognition and resolution in unstructured text](#). *Remote Sensing*, 12(18).
- Takashi Awamura, Daisuke Kawahara, Eiji Aramaki, Tomohide Shibata, and Sadao Kurohashi. 2015. [Location name disambiguation exploiting spatial proximity and temporal consistency](#). In *Proceedings of the 3rd International Workshop on Natural Language Processing for Social Media*, pages 1–9.
- Grant DeLozier, Jason Baldrige, and Loretta London. 2015. [Gazetteer-independent toponym resolution using geographic word profiles](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29(1).
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *5th International Conference on Learning Representations*.
- Milan Gritta, Mohammad Taher Pilehvar, and Nigel Collier. 2018a. [Which Melbourne? augmenting geocoding with maps](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1285–1296.
- Milan Gritta, Mohammad Taher Pilehvar, Nut Limsoopatham, and Nigel Collier. 2018b. [What's missing in geographical parsing?](#) *Language Resources and Evaluation*, 52(2):603–623.
- Sayali Kulkarni, Shailee Jain, Mohammad Javad Hosseini, Jason Baldrige, Eugene Ie, and Li Zhang. 2021. [Multi-level gazetteer-free geocoding](#). In *Proceedings of Second International Combined Workshop on Spatial Language Understanding and Grounded Communication for Robotics*, pages 79–88.
- Michael D. Lieberman, Hanan Samet, and Jagan Sankaranarayanan. 2010. [Geotagging with local lexicons to build indexes for textually-specified spatial data](#). In *Proceedings of 2010 IEEE 26th International Conference on Data Engineering*, pages 201–212.
- Xiao Ling and Daniel Weld. 2012. [Fine-grained entity recognition](#). In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 1, pages 94–100.

- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. [Universal Dependency annotation for multilingual parsing](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, volume 2, pages 92–97.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A Python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Michael Speriosu and Jason Baldridge. 2013. [Text-driven toponym resolution using indirect supervision](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1466–1476.
- Fan Yang and Paul Vozila. 2014. [Semi-supervised chinese word segmentation using partial-label learning with conditional random fields](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 90–98.