

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Geoparsing: Solved or Biased? An Evaluation of Geographic Biases in Geoparsing

Permalink

<https://escholarship.org/uc/item/03z5m5k8>

Author

Liu, Zilong

Publication Date

2022

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Geoparsing: Solved or Biased? An Evaluation of Geographic Biases in Geoparsing

A Thesis submitted in partial satisfaction
of the requirements for the degree

Master of Arts
in
Geography

by

Zilong Liu

Committee in charge:

Professor Krzysztof Janowicz, Chair
Professor Konstadinos Goulias
Professor Kelly Caylor

March 2022

The Thesis of Zilong Liu is approved.

Professor Konstadinos Goulias

Professor Kelly Caylor

Professor Krzysztof Janowicz, Committee Chair

March 2022

Geoparsing: Solved or Biased? An Evaluation of Geographic Biases in Geoparsing

Copyright © 2022

by

Zilong Liu

Acknowledgements

I would like to thank the members of my committee, Professor Konstadinos Goulias and Professor Kelly Caylor, for their guidance on this work and their reviews of my thesis. I would like to thank my advisor, Professor Krzysztof Janowicz, for offering not only the original idea of this thesis but also continuous support and exceptional inputs throughout this work. Jano has always been sharing invaluable insights, giving me the freedom to think outside the box, and most importantly, serving as a real-life example of promoting moral values in this disturbing time. He has led me to a path that I could not foresee before we first met in 2019. I would also like to thank all other members of the STKO Lab, including Rui Zhu, Ling Cai, Meilin Shi, and Gengchen Mai. They have shown me the spirit of staying strong and optimistic about our dreams as well as the importance of thinking critically about our research. All of them have also greatly contributed to my academic growth and beyond. In addition, I would like to thank all my other friends, especially those that I have met during my undergraduate study at East China Normal University and my graduate study at University of California, Santa Barbara. I would like to give special thanks to Jiaqi Liang who has shown me that physical distancing cannot separate humans apart as long as our minds and hearts are intertwined together.

Lastly, I would like to sincerely thank my parents and my grandmother for their unconditional love and support to let me walk freely across the land, roam the ocean, and fly high in the sky. They are every reason why I can shine in this world.

Abstract

Geoparsing: Solved or Biased? An Evaluation of Geographic Biases in Geoparsing

by

Zilong Liu

Geoparsing, the task of extracting toponyms from texts and associating them with geographic locations, has witnessed remarkable progress over the past years. However, despite its intrinsically *geospatial* nature, existing evaluations tend to focus on overall performance while paying little attention to its variation across geographic space. In this work, we attempt to answer the question whether *geoparsing is solved or biased* by conducting a spatially-explicit evaluation, namely an evaluation of the regional variability in geoparsing performance. Particularly, we will analyze the spatial autocorrelation underlying this regional variability. By performing hot and cold spot detection over results of several open-source geoparsers, we observe that none of them performs equally well across geographic space, and some are geographically biased towards some regions but against others. We also carry out a comparative experiment showing that state-of-the-art geoparsers developed with neural networks do not necessarily outperform the off-the-shelf tools across geographic space. To understand the implications behind this observed regional variability, we evaluate geographic biases involved in geoparsing research centered around data contribution and usage, algorithm design, and performance evaluations. Particularly, our spatially-explicit performance evaluation serves as an approach to evaluation bias mitigation in geoparsing. We conclude that previous performance evaluations published in the literature are overly optimistic, thus hiding the fact that geoparsing is far from solved, and geoparsers require debiasing in addition to further considerations when being applied to (geospatial) downstream tasks.

Contents

Abstract	v
1 Introduction	1
2 Related Work	5
2.1 Spatially-Explicit Performance Evaluations	5
2.2 Geoparsing Evaluation	6
3 A Spatially-Explicit Geoparsing Performance Evaluation	9
3.1 Geoparser and Evaluation Corpora	9
3.2 Evaluation Metrics	13
3.3 Spatial Autocorrelation Detection	15
4 Evaluation Results	16
4.1 The Normalized Frequency Distributions of Geoparsing Performance Indicators	16
4.2 Spatial Autocorrelation in Toponym Recognition Performance	21
4.3 Spatial Autocorrelation in Toponym Resolution Performance	21
5 Geographic Biases in Datasets, Algorithms, and Performance Evaluations	29
5.1 Representation Bias	29
5.2 Aggregation Bias	32
5.3 Algorithmic Bias	33
5.4 Evaluation Bias	34
6 Conclusions and Future Work	35
6.1 Conclusions	35
6.2 Future Work	36
Bibliography	39

Chapter 1

Introduction

Geoparsing is a key part of geographic information retrieval [1]. Two consecutive steps constitute the pipeline of geoparsing systems (i.e., geoparsers), namely toponym recognition and toponym resolution. Toponym recognition is often treated as a sub-task of Named Entity Recognition (NER), or more precisely speaking, Named Entity Recognition and Classification (NERC) [2]. It refers to the process of identifying toponyms from texts in various forms of literature, such as news and Wikipedia articles. Each recognized toponym is then fed into a toponym resolution model that selects the correct place reference along with its geo-location information from all candidates with the same place name. Toponym resolution is also referred to as place name disambiguation [3, 4] in the related literature.

In recent years, neural network architectures have contributed to the improvement in the quality of geoparsers. Examples in toponym recognition include Wang et al. [5] that used an improved version of the bidirectional Long Short-Term Memory (LSTM) neural network with a Conditional Random Field (CRF) layer (BiLSTM-CRF) [6], and Hu et al. [7] that used C-LSTM [8] combining a Convolutional Neural Network (CNN) with LSTM. Similarly, Gritta et al. [9] and Kulkarni et al. [10] applied CNN while Fize

et al. [11] used LSTM in toponym resolution. As the use of deep learning techniques has greatly improved geoparsing performance, Wang and Hu [12] recently asked whether the geoparsing performance obtained with state-of-the-art geoparsers is good enough to essentially call the problem *solved*. The authors argued that geoparsing can be claimed to be solved when it comes to prominent place names in well-formatted texts because their evaluation shows that these geoparsers, particularly deep-learning-based ones, can reach an outstanding level of toponym recognition performance and also relatively low errors in toponym resolution on multiple benchmark datasets.

However, there is a loophole in how the performance evaluation is carried out. Current evaluations only measure the overall performance of a geoparser on a corpus, thereby ignoring the fact that, as a *geospatial* task geoparsing also needs to be evaluated from a geospatial perspective. Put differently, how geoparsers perform across geographic space is not evaluated, thus making such evaluations susceptible to biases, e.g., uneven geographic coverage. Recent work has shown spatial heterogeneity in the performance of a deep-learning-based toponym resolution model developed by Fize et al. [11], which fails to work in most of the southern and western regions in the US except several major cities, and similar poor performance is also obtained in small regions in the southwest of France and the north of Japan. Our work tries to expand the scale of geoparsing performance evaluation to both toponym recognition and resolution.

In addition, we aim to look beyond spatial heterogeneity, and study spatial autocorrelation [13, 14, 15] underlying regional variability in geoparsing performance. There are two reasons why we choose spatial autocorrelation as the focus of our performance evaluation. First, spatial autocorrelation is inherent in data with a spatial structure [16, 17, 18], and geoparsing is affected by such autocorrelation effects in place names that also exhibit distance decay patterns in their collective similarity [19]. Second, we consider that an analysis of second-order spatial variations in geoparsing performance will

shed light on where a geoparser is biased towards or against, which we argue is necessary for a spatially-explicit evaluation on geoparsing performance. The term *spatially-explicit* has been frequently used in ecological studies [20, 21, 22], and has been recently regarded as an vital characteristic that a GeoAI [23] should demonstrate by satisfying at least one of four tests, including the invariance test, representation test, formulation test, and outcome test [24]. Examples of such spatially-explicit machine learning models include geo-aware image classification [25] and multi-scale spatial representation learning [26]. Similarly, we consider an evaluation to be spatially-explicit if it fulfills any of the four tests above.

If we observe strong spatial effects in geoparsing performance, this would be an important indicator of geographic biases in existing systems and their evaluations. Geographic biases cause disparities in the geographic distributions between sampled and ground-truth data [27, 28, 29], and they can also lead to quality issues in data contribution to volunteered geographic information through crowdsourcing [30, 31]. Such biases have inspired an interest in sampling bias mitigation [29, 32] and representativeness assessment [33]. Meanwhile, they have also drawn attention from the machine learning community as machine learning researchers have faced with the same lack of geographic diversity in open datasets, such as ImageNet [34] and Open Images [35], which results in biasing their image classifiers towards Europe and North America [36]. In addition, geographic biases also affect spatial data aggregation, which can result in reduced reliability of multivariate statistical analysis [37] and perturbations in feature embeddings that destabilize neural networks used in scenarios such as deep-learning-based traffic predictive modeling [38]. These potential consequences of geographic biases motivate us to investigate geographic bias issues that have not been discussed in geoparsing by studying the datasets and beyond, i.e., algorithms and performance evaluations. Put differently, we examine whether the claim that geoparsing is essentially solved is true, or whether the data, models, and

the usage of evaluation metrics are simply biased.

Our research contributions are as follows:

- Rather than focusing on overall geoparsing performance, we conduct a spatially-explicit evaluation on how geoparsers perform across geographic space. We unveil spatial autocorrelation underlying regional variability in geoparsing performance, and analyze its comparison between deep-learning-based models and off-the-shelf tools in terms of toponym recognition and toponym resolution, respectively.
- We analyze and summarize representation biases, aggregation biases, algorithmic biases, and evaluation biases in geoparsing, along with recent work that attempts to mitigate them. Particularly, our spatially-explicit performance evaluation serves as an approach to evaluation bias mitigation. To the best of our knowledge, our work is the first to provide such geographic bias evaluation in the field of geoparsing (evaluation).

The remainder of this thesis is organized as follows. Chapter 2 provides an overview of related work on spatially-explicit performance evaluations and different kinds of geoparsing evaluation studies. Chapter 3 introduces our spatially-explicit evaluation of geoparsing performance. Chapter 4 describes an exploratory analysis on normalized frequency distributions of geoparsing performance indicators, and the results about our performance evaluation. Chapter 5 discusses geographic biases involved in geoparsing research, and how recent work attempts to mitigate these issues. Finally, we summarize our work and propose future directions in Chapter 6.

Chapter 2

Related Work

2.1 Spatially-Explicit Performance Evaluations

There are two potential avenues for a performance evaluation to become spatially-explicit. First, a spatially-explicit performance evaluation might use evaluation metrics where spatial information (e.g., distance) is incorporated. For instance, Xu and Zhang [39] conducted a sensitivity analysis on land suitability evaluation (LSE), in which the Earth Mover’s Distance is applied to identify spatial variations between the original map and the simulated LSE map. Second, there is usually a geospatial perspective from which a spatially-explicit performance evaluation is carried out. Examples include a “policyscape analysis” for biodiversity conservation [40], an evaluation framework for integrated carbon sequestration and biodiversity conservation [41], and a land use conflict evaluation approach [42].

According to a recent review of evaluation metrics used in geoparsing [43], toponym recognition performance can be evaluated with metrics such as *Precision*, *Recall*, and *F-Score*. *Precision* measures the percentage of correctly-recognized toponyms among all recognized toponyms, and *Recall* measures the percentage of correctly-recognized

toponyms among all annotated toponyms. *F-Score* is the harmonic mean of *Precision* and *Recall*. No spatial information is involved in the calculation of these metrics, thus hindering the spatial explicitness of performance evaluations of toponym recognition. On the other hand, evaluation metrics are mostly distance-based for toponym resolution. Commonly-used metrics include *Mean Error Distance (MED)*, *Median Error Distance (MdnED)*, *Accuracy@161*, and *Area Under the Curve (AUC)*. These metrics evaluate toponym resolution performance by comparing the error distance determined by how far a resolved location is from its corresponding annotated location. *MED*, *MdnED*, and *AUC* calculate the mean, median, and overall deviation of error distances, respectively. *Accuracy@161* is used to calculate the percentage of correctly-resolved locations among all annotated locations. It considers a distance threshold of 161 kilometers, within which a resolved location is regarded as correct. However, when using these metrics to evaluate toponym resolution, the performance across geographic space was often not considered.

2.2 Geoparsing Evaluation

Geoparsing evaluation is an important part of geoparsing research, which is concerned with concrete metrics, progress reviews, reproducibility issues, and benchmark dataset construction in geoparsing.

Gritta et al. [44] discussed standard metrics used to evaluate geoparsing performance, and provided an evaluation framework. For example, they divided toponym resolution evaluation metrics into coordinate-based, set-based, and ranking-based metrics. Also, the authors highlighted that spatial scopes of geoparsing, i.e., whether a geoparser is applied to a local or global coverage, should be taken into considerations in performance evaluations.

Wang and Hu [43] moved one step forward by not only conducting a more compre-

hensive review of evaluation corpora, state-of-the-art models, and evaluation metrics. The authors built an Extensible and Unified Platform for Evaluating Geoparsers (EUPEG)¹, which is a benchmark platform aimed at improving reproducibility in geoparsing research for comparative experiments. Both geoparsers and evaluation corpora hosted on this platform are open-source. These geoparsers use different toponym recognition techniques, such as general NER tools (e.g., Stanford NER²) and in-house NER tools (e.g., LT-TTT2 [45]). In the meantime, various toponym resolution models were adopted, such as heuristic rule-based models (e.g., CLAVIN³), geostatistical models (e.g., TopoCluster [46]), and deep-learning-based models (e.g., CamCoder [9]). In addition, there are multiple categories of evaluation corpora available, including news articles (e.g., TR-News [47]), Wikipedia articles (e.g., WikToR [48]), social media posts (e.g., GeoCorpora [49]), and web pages (e.g., Hu2014 [50]).

Along with eight annotated benchmark datasets and eight evaluation metrics on top of EUPEG, Wang and Hu [12] carried out a performance evaluation on nine SOTA neural-network-based geoparsers, including those already hosted on the platform and others developed by top-ranked teams in a geoparsing competition called SemEval-2019 Task 12. They discussed the circumstance under which geoparsing can be considered as solved, and introduced three future directions in geoparsing. These directions include population-free toponym resolution, fine-grained geoparsing, and the usage of additional gazetteers in toponym resolution. In this thesis, we will showcase that their performance evaluation would benefit from being more spatially-explicit, because such spatially-explicit evaluations will allow us to compare geoparsing performance among different locations.

More recently, Laparra and Bethard [51] combined Wikipedia⁴ and OpenStreetMap⁵

¹<https://geoai.geog.buffalo.edu/EUPEG/>

²<https://nlp.stanford.edu/software/CRF-NER.html>

³<https://github.com/Novetta/CLAVIN>

⁴<https://www.wikipedia.org/>

⁵<https://www.openstreetmap.org/>

to construct a new kind of benchmark dataset for geoparsing compositions of place mentions into geographic regions. The authors also proposed accompanying evaluation metrics that can be used to compare predicted geometries with ground-truth geometries in either a strict or relaxed way. As our intention is to reproduce studies on geoparsing individual toponyms rather than complex geographic descriptions, we do not use their evaluation framework in our study.

Chapter 3

A Spatially-Explicit Geoparsing Performance Evaluation

3.1 Geoparser and Evaluation Corpora

Previous work has studied toponym recognition and resolution individually, and therefore, we evaluate each of them separately. The criterion of our model selection is that their relevant resources, including evaluation corpora (and training corpora, if applied), are available online so that they can be easily reproduced as baseline models for geoparsing research. For both toponym recognition and resolution, we use a deep-learning-based model that has achieved state-of-the-art results and an off-the-shelf tool to analyze whether the former necessarily outperforms the latter across geographic space.

Toponym Recognition Models For toponym recognition, we choose a pre-trained version of NeuroTPR¹ and the named entity recognition module of spaCy² (version 2.1)

¹<https://github.com/geoai-lab/NeuroTPR>

²<https://spacy.io>

coupled with a large-sized trained English pipeline³. NeuroTPR is a BiLSTM-conditional random field toponym recognition model that deals with social media messages [5], and spaCy is an open-source Python library for natural language processing.

Toponym Resolution Models For toponym resolution, we compare a pre-trained version of CamCoder⁴, a CNN-based toponym resolution model that integrates both lexical and geographic knowledge [9], with the rule-based Edinburgh Geoparser⁵ [52] (version 1.2).

Toponym Recognition Corpus To evaluate toponym recognition, we select GeoCorpora [49], a social media corpus containing geo-annotated tweets along with place mentions (see their geographic distribution in Figure 3.1). This dataset was also used as one of the two evaluation corpora in Wang et al. [5], and is available on Github⁶. Among all annotated toponyms in the 2,122 tweets, only those that have labeled coordinates are covered in our evaluation.

Toponym Resolution Corpora For toponym resolution evaluation, we select LGL [53], GeoVirus [9], and WikToR [48], which were also used in the evaluation of Gritta et al. [9] and shared along with CamCoder⁴. Geographic distributions of annotated locations in LGL, GeoVirus, and WikToR can be seen from Figure 3.2, Figure 3.3, and Figure 3.4, respectively. Both LGL and GeoVirus are news corpora. LGL is also the most frequently-used geoparsing evaluation dataset, consisting of 588 news articles from 78 local newspapers. GeoVirus contains 229 news articles from WikiNews, and they are centered around global disease outbreaks and epidemics. WikToR is a Wikipedia dataset

³https://spacy.io/models/en#en_core_web_lg

⁴<https://github.com/milangritta/Geocoding-with-Map-Vector>

⁵<https://www.ltg.ed.ac.uk/software/geoparser/>

⁶<https://github.com/geovista/GeoCorpora>

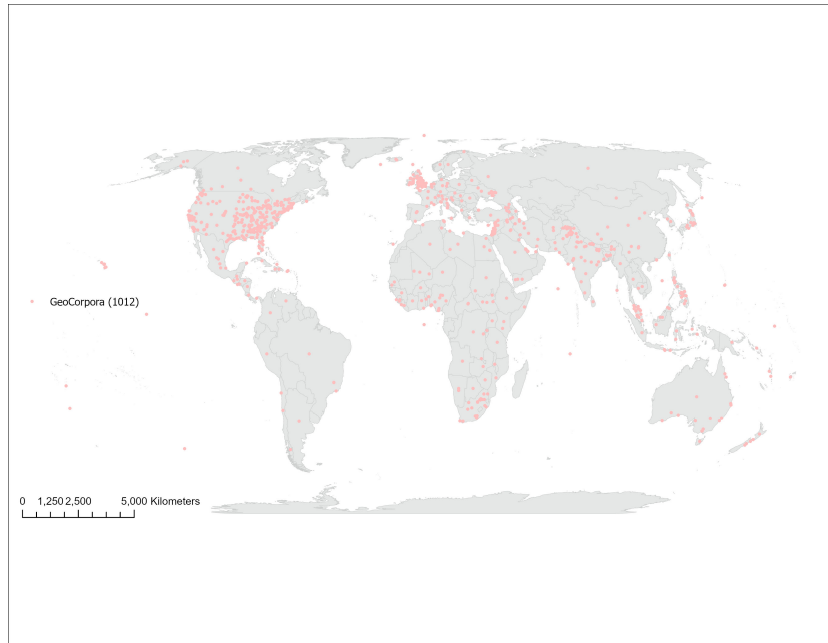


Figure 3.1: The geographic distribution of annotated locations in GeoCorpora

containing 5,000 articles within which annotated toponyms are widely distributed across the world. The reason why we choose the three corpora here is that we want to analyze how CamCoder performs across geographic space on benchmark datasets with different levels of ambiguity. According to Gritta et al. [9], WikToR has higher place name ambiguity than GeoVirus, and GeoVirus has higher ambiguity than LGL.

Kulkarni et al. [10] found that WikToR contains wrong coordinates for some places because the sign of either their latitude or longitude is flipped, and in all three corpora the same toponym will have slightly different coordinates because they were created differently. To correct these inconsistencies, we follow their method to unify three corpora by using the shared data patches⁷. Annotated toponyms without location information in the data patches are not included in our study.

⁷https://github.com/google-research-datasets/mlg_evaldata

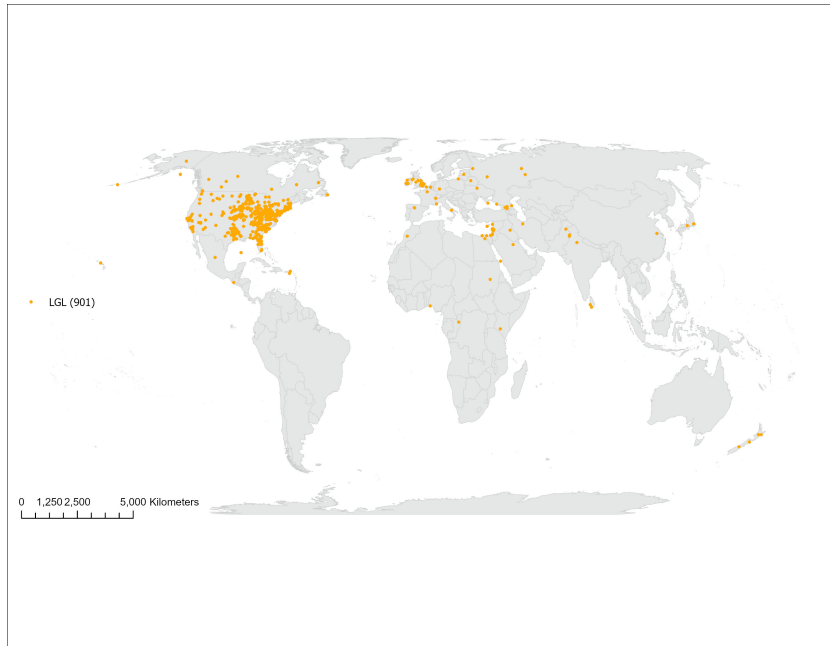


Figure 3.2: The geographic distribution of annotated locations in LGL

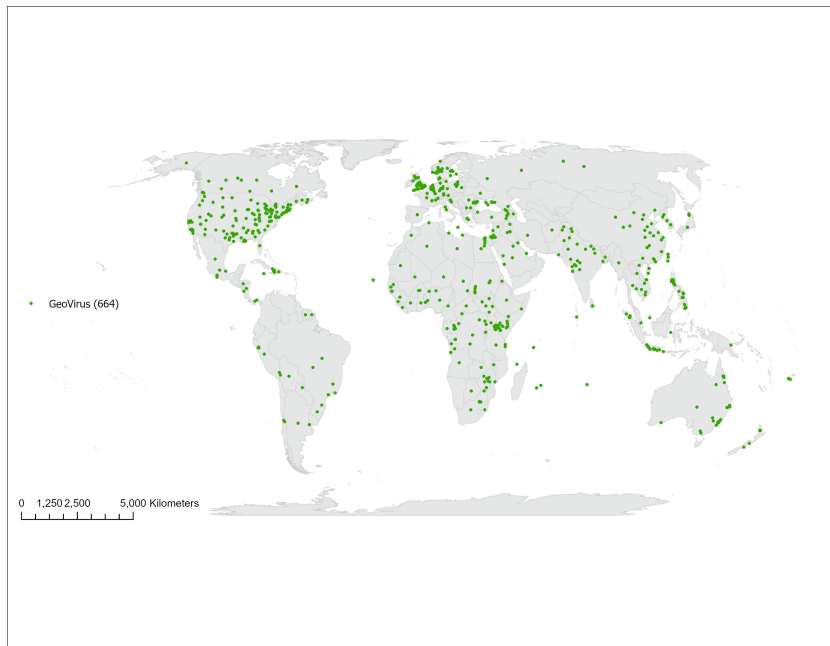


Figure 3.3: The geographic distribution of annotated locations in GeoVirus

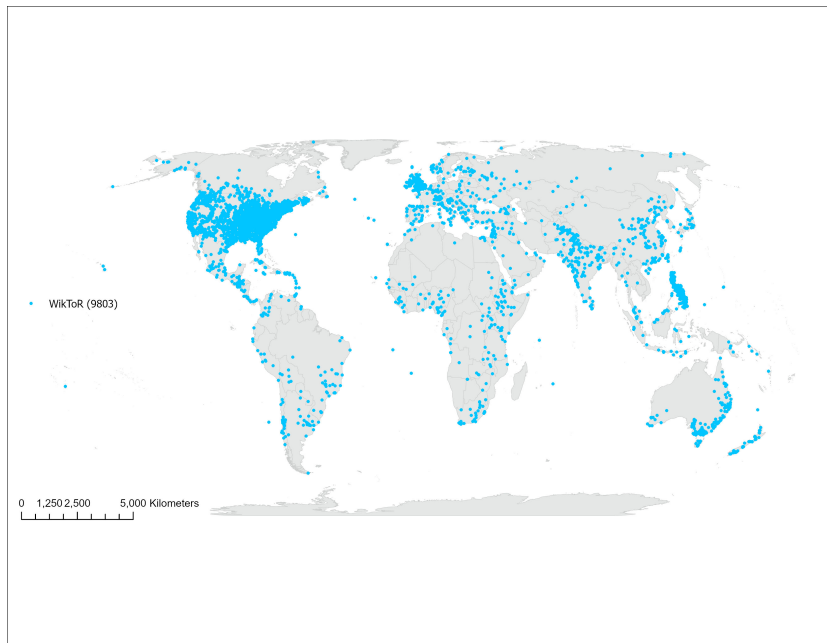


Figure 3.4: The geographic distribution of annotated locations in WikToR

3.2 Evaluation Metrics

In our performance evaluation, we select *Recall* and *MdnED* as they are frequently used to evaluate the performance of toponym recognition and toponym resolution, respectively. We apply them to measuring geoparsing performance at each annotated location individually, so that we can disclose how geoparsers perform across geographic space.

Toponym Recognition Evaluation Metric For toponym recognition, we use *Recall* to measure the proportion of the number of times an annotated location being correctly identified by a model among the number of times it is annotated (in a corpus). In Equation 3.1, $Recall_i$ is the *Recall* of the i^{th} annotated location in a corpus; tp_i is the number of times it is recognized; and fn_i is the number of times it fails to be recognized. The range of *Recall* is $[0, 1]$. A higher $Recall_i$ indicates a better toponym recognition performance with respect to the i^{th} annotated location. We adopt exact matching [48],

meaning that only toponyms that match exactly with their ground-truth annotations are considered valid.

$$Recall_i = \frac{tp_i}{tp_i + fn_i} \quad (3.1)$$

Toponym Resolution Evaluation Metric For toponym resolution, *MdnED* is a coordinate-based metric calculated as the median of error distances from the location of an annotated toponym to its resolved location (in a corpus). Compared with the commonly-used *MED* that calculates the mean, *MdnED* is better at dealing with outliers that exist in computed error distances, and therefore, *MdnED* can minimize the distortion of evaluation results. In Equation 3.2, $MdnED_i$ is the *MdnED* of the i^{th} annotated location in a corpus; ed_{ij} is its j^{th} error distance computed; $\vec{x}_i = (x_i, y_i)$ is its annotated location; n_i is the number of its resolved locations; $\vec{x}_{ij} = (x_{ij}, y_{ij})$ is its j^{th} resolved location; and $Dist(\cdot, \cdot)$ is the error distance between a pair of geographic coordinates, computed as the great circle distance in our experiment. A larger *MdnED* indicates worse toponym resolution performance with respect to an annotated location. Only the annotated locations that are recognized by a geoparser and can be found in the GeoNames gazetteer⁸ are considered in the computation of *MdnED*. A shared version of GeoNames by Gritta [54] is used to reproduce CamCoder, while the online version of GeoNames is directly accessed by Edinburgh Geoparser during place name disambiguation.

$$MdnED_i = Median(\{ed_{ij} | ed_{ij} = Dist(\vec{x}_i, \vec{x}_{ij}), \forall j \in [1, n_i]\}) \quad (3.2)$$

⁸<https://www.geonames.org/>

3.3 Spatial Autocorrelation Detection

To evaluate spatial autocorrelation effects in geoparsing performance, we apply the Getis-Ord G_i^* statistic [55, 56] that can help visually reveal hot spots and cold spots in geoparsing performance. In Equation 3.3, G_i^* is the Getis-Ord G_i^* statistic of the i^{th} annotated location (in a corpus); v_i and v_j denote the geoparsing performance indicators (i.e., *Recall* or *MdnED*) of the i^{th} and the j^{th} annotated locations, respectively; \bar{V} is the average geoparsing performance indicator of all annotated toponyms, and N is the number of all annotated locations; w_{ij} is the spatial weight between the i^{th} and the j^{th} locations. We assign 1 to w_{ij} if the j^{th} annotated location is in the neighborhood of the i^{th} annotated location, and 0 otherwise. The Getis-Ord G_i^* statistic is a z-score. A larger positive (or negative) z-score indicates a stronger clustering effect of high (or low) values, which represents hot (or cold) spots. We select the K-nearest neighbors of an annotated location as its neighborhood when calculating the Getis-Ord G_i^* statistics. K is defined as 8 in the experiment.

$$G_i^* = \frac{\sum_{j=1}^N w_{ij}v_j - \bar{V} \sum_{j=1}^N w_{ij}}{\sqrt{\frac{\sum_{j=1}^N v_j^2}{N} - (\bar{V})^2} \sqrt{\frac{N \sum_{j=1}^N w_{ij}^2 - (\sum_{j=1}^N w_{ij})^2}{N-1}}} \quad (3.3)$$

While focusing on how geoparsers perform across space in general, we are also interested in whether regional variability in toponym resolution performance exists for highly ambiguous toponyms as toponym resolution is sensitive to ambiguity determined by the frequency of places with the same name. Therefore, we calculate the standard deviations of *MdnED* produced by both CamCoder and Edinburgh Geoparser for highly ambiguous place names in WikToR, since it has the highest ambiguity among all three evaluation corpora for toponym resolution in our experiment. The standard deviations are reported in Chapter 4.3.

Chapter 4

Evaluation Results

4.1 The Normalized Frequency Distributions of Geoparsing Performance Indicators

First, we provide an exploratory analysis showing that the performance indicator distributions of both toponym recognition and toponym resolution are highly skewed. Figure 4.1 and Figure 4.2 describe the normalized frequency distributions of *Recall* for spaCy and NeuroTPR, respectively. We can see there is a peak indicating that more than 50% annotated locations have a *Recall* ranging from 0 to 0.1 for both toponym recognition models. Figure 4.3, Figure 4.4, and Figure 4.5 show the normalized frequency distributions of *MdnED* for Edinburgh Geoparser with respect to LGL, GeoVirus, and WikToR, respectively, while Figure 4.6, Figure 4.7, and Figure 4.8 show those distributions for CamCoder. On the contrary to *Recall*, *MdnED* is expected to be as small as possible, but in most cases it can reach up to 16,000 kilometers.

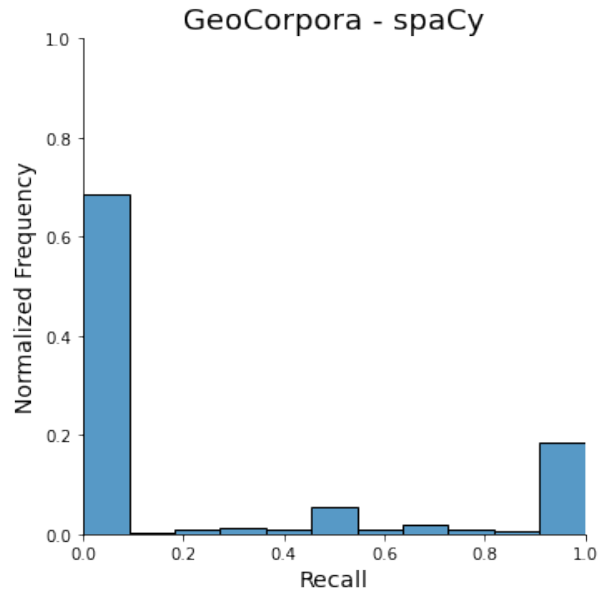


Figure 4.1: The normalized frequency distributions of *Recall* with respect to all annotated locations in GeoCorpora for spaCy

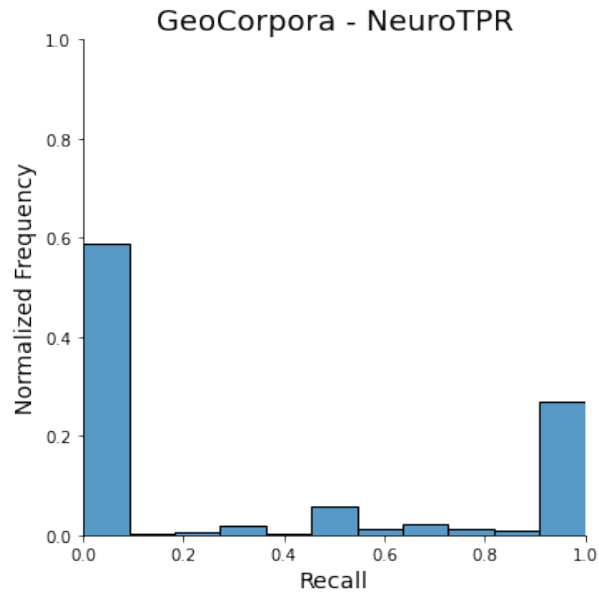


Figure 4.2: The normalized frequency distributions of *Recall* with respect to all annotated locations in GeoCorpora for NeuroTPR

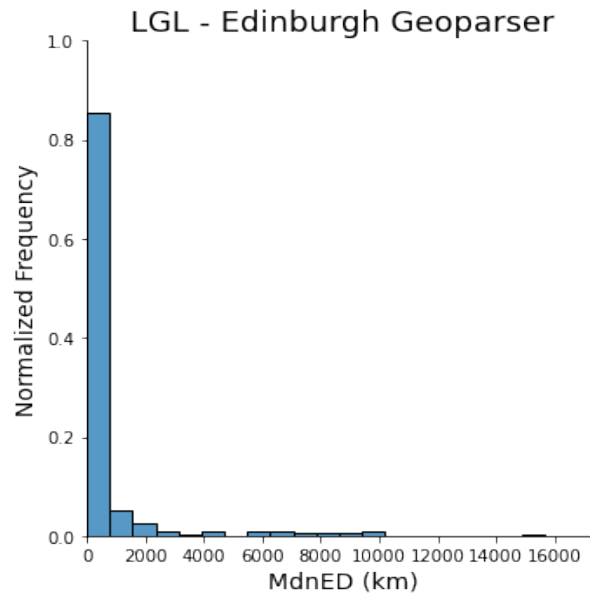


Figure 4.3: The normalized frequency distributions of $MdnED$ with respect to all annotated locations in LGL for Edinburgh Geoparser




Figure 4.4: The normalized frequency distributions of $MdnED$ with respect to all annotated locations in GeoVirus for Edinburgh Geoparser

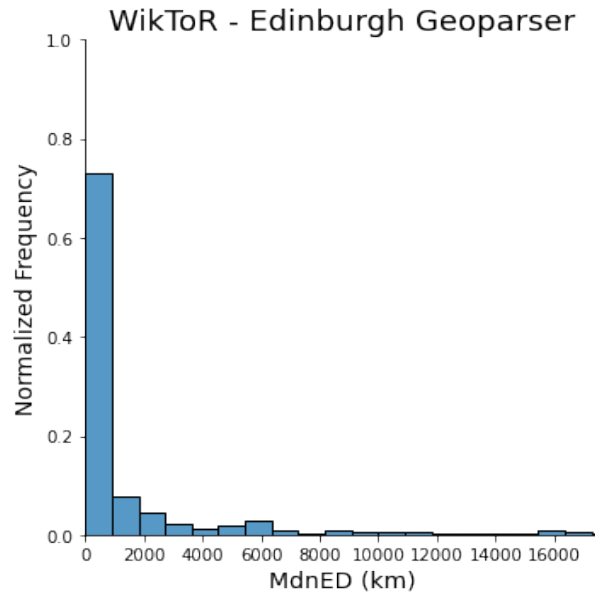


Figure 4.5: The normalized frequency distributions of $MdnED$ with respect to all annotated locations in WikToR for Edinburgh Geoparser

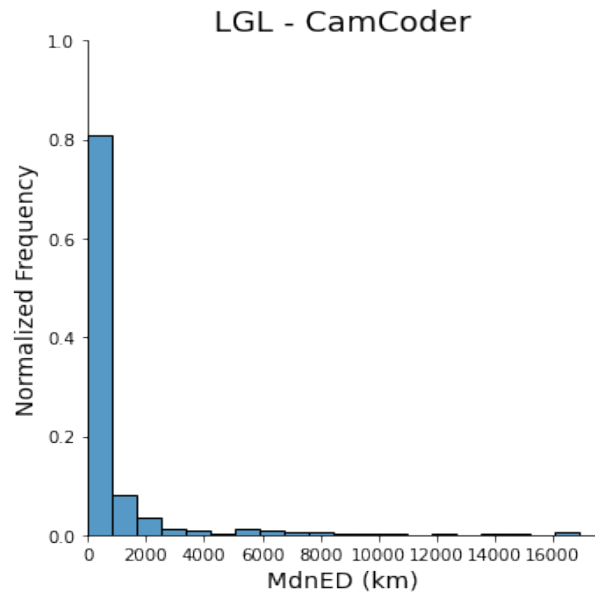


Figure 4.6: The normalized frequency distributions of $MdnED$ with respect to all annotated locations in LGL for CamCoder

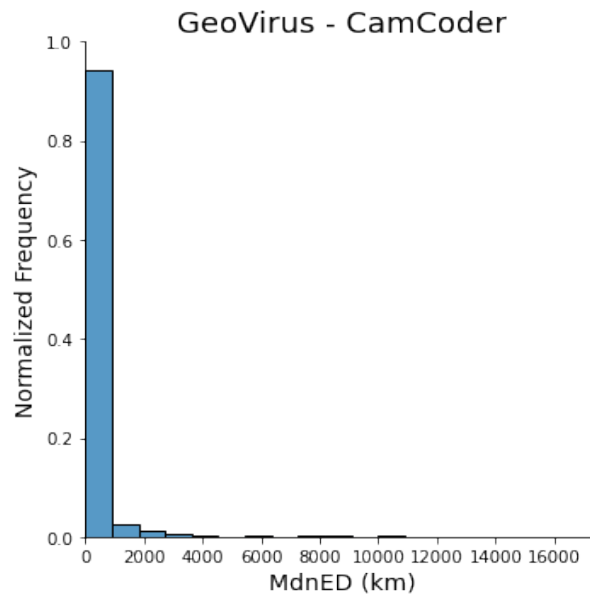


Figure 4.7: The normalized frequency distributions of $MdnED$ with respect to all annotated locations in GeoVirus for CamCoder

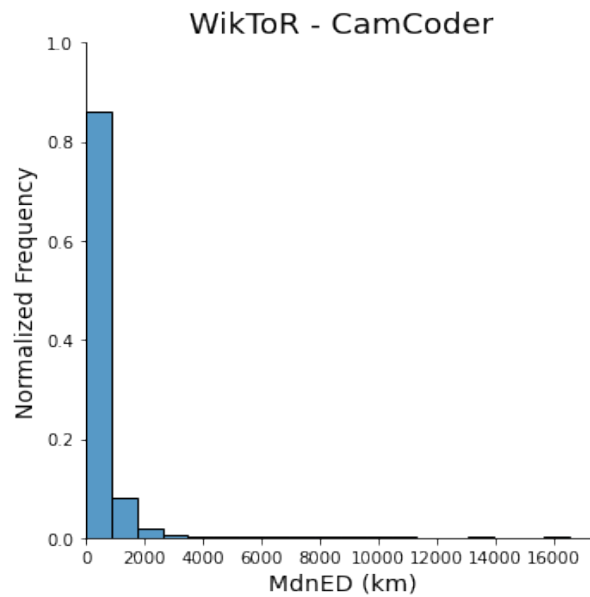


Figure 4.8: The normalized frequency distributions of $MdnED$ with respect to all annotated locations in WikToR for CamCoder

4.2 Spatial Autocorrelation in Toponym Recognition Performance

Figure 4.9 shows that both hot spots and cold spots are detected in toponym recognition performance of spaCy on GeoCorpora. For spaCy, most of all the hot spots are found in the United States, the United Kingdom, and Malaysia on a global scale. Cold spots can be seen not only in several South American countries near the equator, Middle East countries (e.g., Bahrain and Qatar), South Asia countries (e.g., India), and so forth, but also in the United States and the United Kingdom, where hot spots have been found as well. Figure 4.10 shows hot spots and cold spots in toponym recognition performance of NeuroTPR on GeoCorpora, where there is a larger proportion of more cold spots with 95% confidence compared with Figure 4.9. Again, we can see coexistence of both hot spots and cold spots in the United States and the United Kingdom.

In general, the observed regional variability in toponym recognition performance shows that toponym recognition is geographically biased towards some regions and against others, on both global and local scales.

4.3 Spatial Autocorrelation in Toponym Resolution Performance

In terms of toponym resolution, however, only cold spots can be found in the performance of both Edinburgh Geoparser and CamCoder on LGL, GeoVirus, and WikToR. Figure 4.11, Figure 4.13, and Figure 4.15 show that the geographic distribution of cold spots covers more regions on a global scale as the place name ambiguity of the corpus increases from LGL, GeoVirus to WikToR. Cold spots are mostly observed in the United

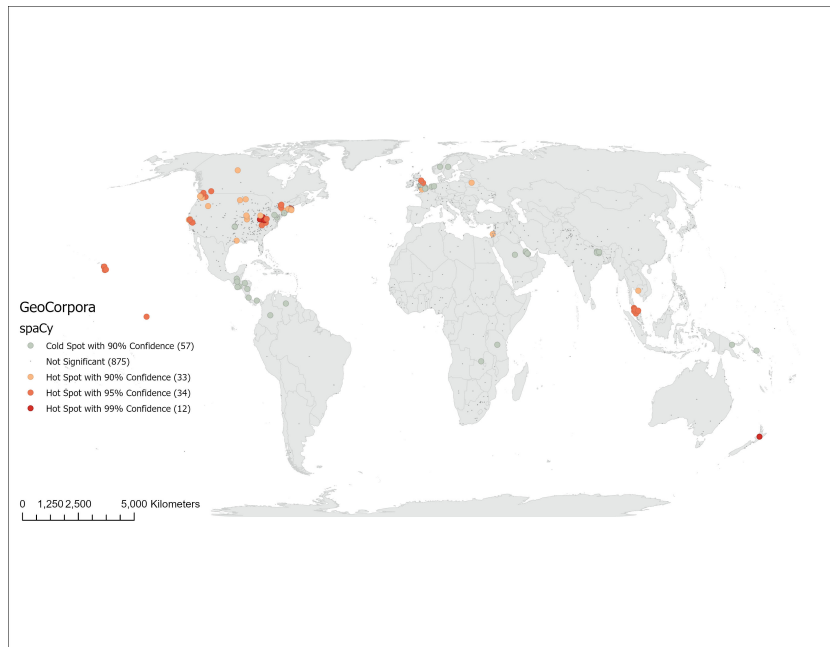


Figure 4.9: Hot spots and cold spots in toponym recognition performance of spaCy on GeoCorpora

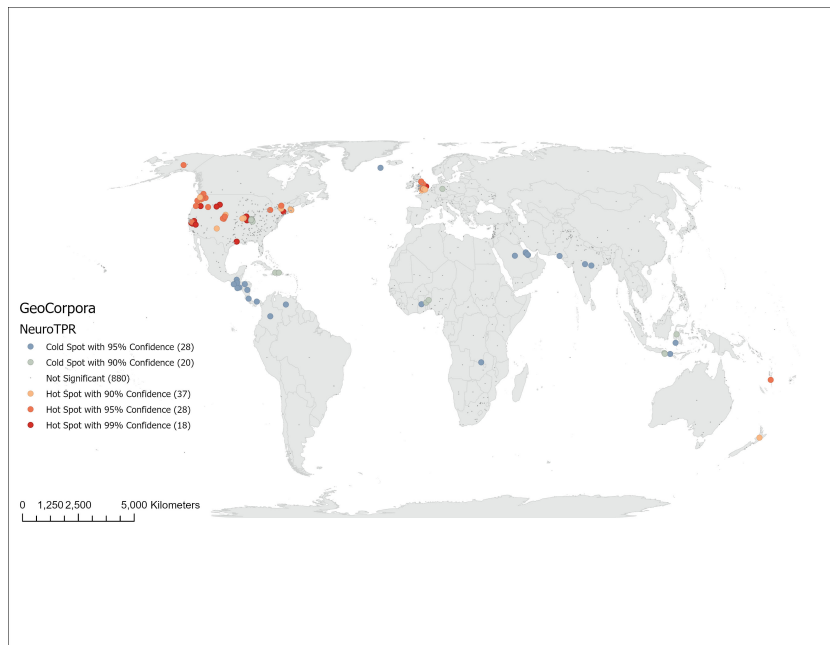


Figure 4.10: Hot spots and cold spots in toponym recognition performance of NeuroTPR on GeoCorpora

States for LGL, in the United States, Mexico, the United Kingdom, Australia, and Fiji for GeoVirus, and in the United States, Canada, the Philippines, Australia, and New Zealand for WikToR. Similar observations can be found for CamCoder’s toponym resolution performance on LGL, GeoVirus, and WikToR, which are shown in Figure 4.12, Figure 4.14, and Figure 4.16, respectively.

In addition, CamCoder produces more cold spots (with higher confidence), and these cold spots are widely distributed across the world. By comparing the toponym resolution performance between Edinburgh Geoparser and CamCoder, we can see while the overall performance of CamCoder is higher than Edinburgh Geoparser as reported in previous research [9, 12], there are more regions (e.g., South America and Africa) where it fails to perform well. Note that the total of cold spots produced by CamCoder is greater than that of Edinburgh Geoparser because these two geoparsers use different toponym recognition models. Put aside this difference, the wider geographic distribution of cold spots in toponym resolution performance of CamCoder indicates a stronger geographic bias in this deep-learning-based model in comparison with the rule-based Edinburgh Geoparser.

In general, the observed regional variability in toponym resolution performance shows that toponym resolution is even worse than toponym recognition because there is no hot spots found. Put differently, toponym resolution is simply biased against many regions, both globally and locally.

The Standard Deviations of *MdnED* for Highly Ambiguous Toponyms in

WikToR Then, we examine the top ten most ambiguous place names in WikToR, and analyze regional variability in their corresponding toponym resolution performance according to the standard deviation of *MdnED*, as shown in Table 4.1. We find that Edinburgh Geoparser fails to resolve half of the selected toponyms, which is because



Figure 4.11: Cold spots in toponym resolution performance of Edinburgh Geoparser on LGL

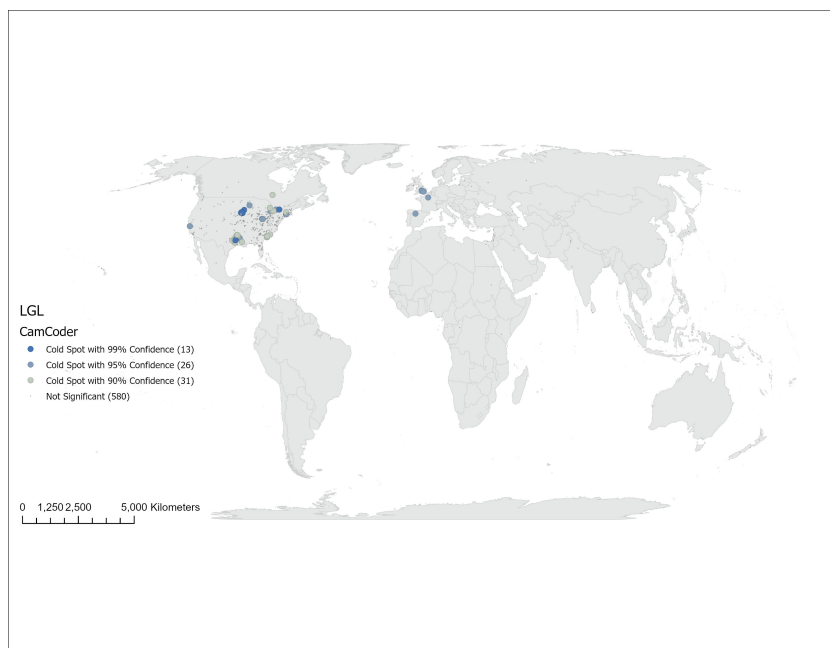


Figure 4.12: Cold spots in toponym resolution performance of CamCoder on LGL



Figure 4.13: Cold spots in toponym resolution performance of Edinburgh Geoparser on GeoVirus

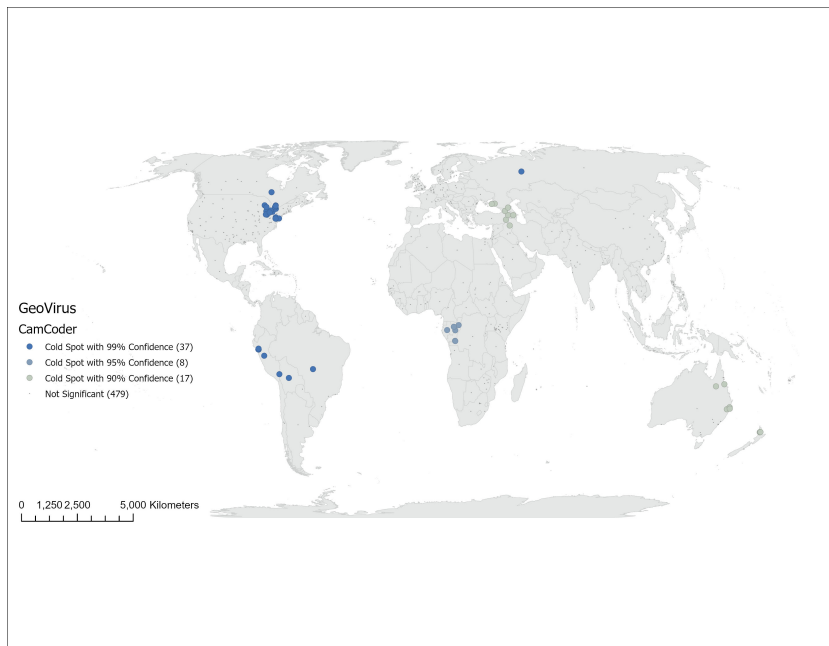


Figure 4.14: Cold spots in toponym resolution performance of CamCoder on GeoVirus

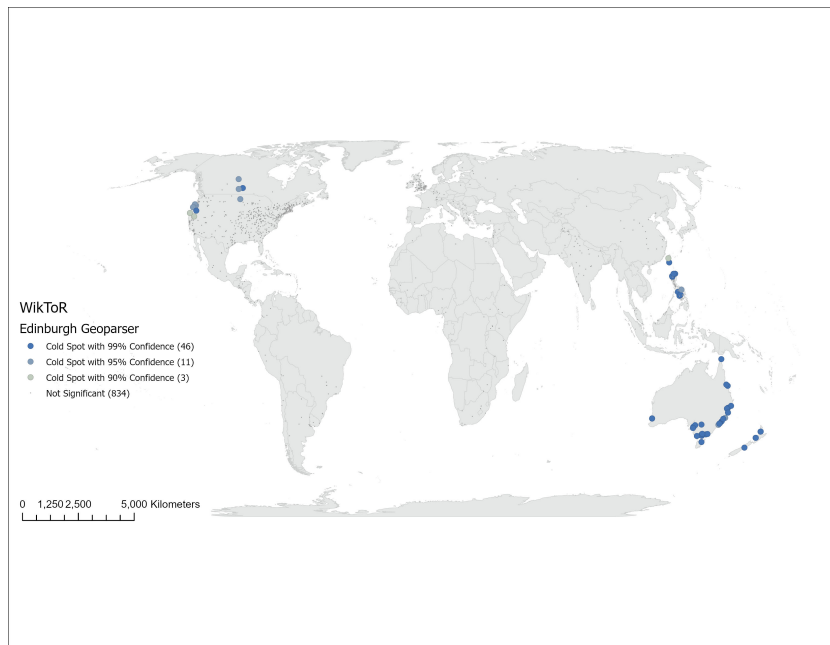


Figure 4.15: Cold spots in toponym resolution performance of Edinburgh Geoparser on WikToR

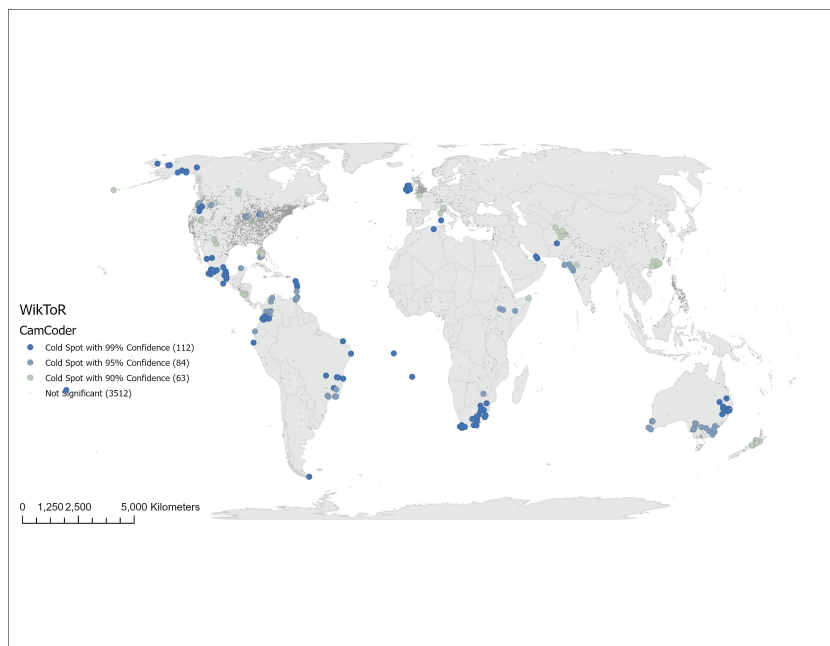


Figure 4.16: Cold spots in toponym resolution performance of CamCoder on WikToR

of a failure either to identify them during toponym recognition or to retrieve their coordinates from the gazetteer. For toponyms resolved only by CamCoder, we observe very large standard deviations of *MdnED* for Springfield and Georgetown, which are 2,770.63 and 2,742.79 kilometers, respectively. For toponyms (including Washington County, Greenville, Kingston, Hamilton, and Newport) resolved by both models, Edinburgh Geoparser results in greater standard deviations of *MdnED* than CamCoder. Particularly, this difference is extremely large for Hamilton and Newport. As all these standard deviations (except the one for Newport resolved by CamCoder) are greater than 161 kilometers, which is regarded as a threshold within which a resolved location can be considered correct relative to its annotated location [43], both models exhibit very strong regional variability in toponym resolution performance for place names with high ambiguity.

Table 4.1: The standard deviations (SD) of *MdnED* for highly ambiguous toponyms in WikToR

Toponym	Model	MdnED SD (km)
Washington County	Edinburgh Geoparser	807.70
	CamCoder	692.18
Clinton	Edinburgh Geoparser	/
	CamCoder	539.43
Greenville	Edinburgh Geoparser	407.52
	CamCoder	514.68
Springfield	Edinburgh Geoparser	/
	CamCoder	2770.63
Georgetown	Edinburgh Geoparser	/
	CamCoder	2742.79
Kingston	Edinburgh Geoparser	680.48
	CamCoder	226.45
Franklin County	Edinburgh Geoparser	/
	CamCoder	414.99
Hamilton	Edinburgh Geoparser	4105.57
	CamCoder	622.24
Jefferson County	Edinburgh Geoparser	/
	CamCoder	384.38
Newport	Edinburgh Geoparser	5705.19
	CamCoder	145.05

Chapter 5

Geographic Biases in Datasets, Algorithms, and Performance Evaluations

While we have shown that geoparsing is geographically *biased* instead of being *solved*, we are also interested in the implications behind the observed autocorrelation effects in its performance, which is likely to be attributed to geographic biases in datasets, algorithms, and beyond that are involved in geoparsing. In this section, we analyze these biases, and discuss their potential influence on geoparsing. While following bias categorization of the machine learning community [57], we also hope to highlight their geospatial characteristics and to point out that they have been (un)intentionally introduced in geoparsing.

5.1 Representation Bias

The first kind of geographic bias is the representation bias in training/evaluation corpora. Here we provide quantitative measurements of how representation biases in datasets

commonly used in recent geoparsing research are across geographic space. Two metrics, *Spatial Misalignment (SM)* and *Spatial Diversity Misalignment (SDM)*, are used to measure the difference in place coverage and the difference in the geographic diversity of place coverage, respectively, between a corpus and the GeoNames Gazetteer. These metrics are introduced in Quattrone et al. [58] to compare the content mapped by power users, i.e., users that represent only a small portion of the entire OpenStreetMap community but produce most of the content, and the content mapped by the crowd. In our experiment, each country/region is divided into $10\text{km} \times 10\text{km}$ grids, and each grid cell is assigned to the number of places within it (in a corpus or in the gazetteer). The country/region data used is a 1:10m shapefile from Natural Earth 5.0.0¹.

Equation 5.1 describes the calculation of *SM*, in which the i^{th} element in the vector \vec{g}_c (or \vec{g}_d) is the number of places in the i^{th} grid of a country/region mapped by a corpus (or a gazetteer).

$$SM = 1 - \frac{\vec{g}_c \cdot \vec{g}_d}{\|\vec{g}_c\| \cdot \|\vec{g}_d\|} \quad (5.1)$$

Equation 5.2 describes the calculation of *SDM*, which is an extension of the Shannon’s Diversity Index [59]. The variable $g_{c,i}$ (or $g_{d,i}$) is the i^{th} element in \vec{g}_c (or \vec{g}_d). As many regions of the world are sparsely populated with annotated locations in LGL and GeoVirus, we only provide measurements for global-scale evaluation corpora used in our spatially-explicit performance evaluation, i.e., GeoCorpora and WikToR. We also provide the measurements for a training corpus named GeoWiki shared by Gritta [54]. This is a version of the English Wikipedia dump that contains 1.4M Wikipedia articles, and was used in the training process of CamCoder without overlapping with WikToR. The English Wikipedia dump has commonly served as an easily-accessible and frequently-

¹<https://www.naturalearthdata.com/>

updated large corpus for the training purpose of many other toponym resolution models and toponym recognition models as well.

$$SDM = \frac{(-\sum_{g_{c,i} \in \vec{g}_c} g_{c,i} \ln g_{c,i}) - (-\sum_{g_{d,i} \in \vec{g}_d} g_{d,i} \ln g_{d,i})}{\max(-\sum_{g_{c,i} \in \vec{g}_c} g_{c,i} \ln g_{c,i}, -\sum_{g_{d,i} \in \vec{g}_d} g_{d,i} \ln g_{d,i})} \quad (5.2)$$

Table 5.1 and Table 5.2 show SM and SDM for GeoWiki, GeoCorpora, and WikToR, respectively. For all three corpora the median of SM is greater than or equal to 0.49, meaning there is a strong misalignment between their place coverage and place coverage of GeoNames. This spatial misalignment is more prominent for evaluation corpora as both their first quantiles of SM are greater than 0.8. In addition, SDM for all three corpora is generally negative. This points out a strong misaligned geographic diversity in their place coverage, particularly for the social media corpus GeoCorpora with a SDM median of -1.00.

Table 5.1: *Spatial Misalignment* between training/evaluation corpora and the GeoNames gazetteer

Dataset	Genre	Min	1st Qu.	Median	3rd Qu.	Max
GeoWiki	Wikipedia	0.00	0.31	0.49	0.63	0.87
GeoCorpora	Social Media	0.35	0.90	0.96	0.99	1.00
WikToR	Wikipedia	0.49	0.87	0.92	0.96	1.00

Table 5.2: *Spatial Diversity Misalignment* between training/evaluation corpora and the GeoNames gazetteer

Dataset	Genre	Min	1st Qu.	Median	3rd Qu.	Max
GeoWiki	Wikipedia	-1.00	-0.30	-0.22	-0.14	0.43
GeoCorpora	Social Media	-1.00	-1.00	-1.00	-0.86	-0.43
WikToR	Wikipedia	-1.00	-1.00	-0.83	-0.70	-0.26

However, it is worth noting that there are also representation biases involved in gazetteer data that have often been used as references for geoparsers to search for candidate place information of a toponym. McDonough et al. [60] questioned the colonist perspective from which databases such as GeoNames, the Alexandria Digital Library, and Wikipedia are built. They also pointed out the lack of temporal metadata for historical toponyms in gazetteers that have already been found to contain inadequate geographic knowledge about many parts of the world. This raises concerns that biases might accumulate as digital resources would replicate the content of their predecessors during their creation process. Therefore, we consider the study of historical aspects of representation biases involved in place name data as an intriguing future direction, which has now been made possible by tremendous efforts in applying knowledge graphs to building historical gazetteers such as Grossner and Mostern [61].

5.2 Aggregation Bias

The second kind of geographic bias is the aggregation bias involved in toponym resolution. In many recent studies, the Earth’s surface is divided into grid cells, and toponym resolution is approached as a classification task where the model predicts the most likely cell the current toponym should fall into based on loss function minimization. However, different discretization of the space will yield different prediction results. In the evaluation of Kulkarni et al. [10], this Modifiable Areal Unit Problem (MAUP) [62] in the prediction of toponym resolution is found to cause a trade-off between model generalization and prediction quality: finer granularity results in higher accuracy in denser regions with more toponyms in training data, while coarser granularity leads to better generalization over data on both global and local scales.

In addition to the prediction process, different patterns can also be learned by to-

ponym resolution models when choosing different granularity during the training process. Coupled with joint minimization of losses at each level, the multi-level neural network architecture proposed by Kulkarni et al. [10] can be one of the many possible solutions to dealing with the MAUP in the training of toponym resolution models.

5.3 Algorithmic Bias

Besides datasets, there is a common algorithmic bias intentionally introduced in toponym resolution. For instance, toponym resolution tends to prefer places with the largest population during place name disambiguation. While applying this simple population heuristic reduces computational complexity, and even brings about better performance than building more complex architecture in some circumstances, its limitations are evident. First, population information needs to be retrieved from gazetteers, and therefore, toponym resolution is made gazetteer-constrained. Second, population information serves as a geographic bias that hinders the fairness towards places with a smaller population in toponym resolution. Gazetteer-free toponym resolution has been studied as one of the many possible solutions recently, and has shown promising results. Examples include topic modeling for place name disambiguation [4], modeling geographic profiles of words [46], spatial language representation learning at multiple levels [10], and toponym co-occurrence representation learning [11]. It is worth noting that gazetteer-free toponym resolution models can even output toponyms that are not inventoried in the gazetteer in the first place.

5.4 Evaluation Bias

Lastly, there is an evaluation bias [63] caused by how geoparsing performance was measured in previous research. How a geoparser performs across geographic space has not been taken into account, since the evaluation metrics are merely used to provide the overall performance instead of the performance at each annotated location (in a corpus). Our spatially-explicit performance evaluation has addressed this issue. However, the spatial (diversity) misalignment analysis in Chapter 5.1 raises another concern that geoparsing performance evaluation is more than geographically biased for lack of a geospatial perspective, because the evaluation bias in geoparsing can be exacerbated by the representation bias in these benchmark evaluation datasets.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

In this work, we presented a spatially-explicit evaluation of geoparsing performance across geographic space. We utilized *Recall* and *MdnED* to measure toponym recognition and resolution performance, respectively, and compared how deep-learning-based models that were claimed to achieve state-of-the-art results and their off-the-self counterparts perform across geographic space. By visualizing the normalized frequency distributions of the two geoparsing performance indicators, we discovered that all normalized frequency distributions are highly skewed. Then, we analyzed the spatial autocorrelation underlying regional variability in geoparsing performance by calculating the Getis-Ord G_i^* statistic of *Recall* and *MdnED* at all annotated locations. We detected hot spots and cold spots in geoparsing performance, which reveals that geoparsing is geographically biased towards some regions and against others. Particularly, in toponym resolution that is sensitive to place name ambiguity, only cold spots were observed, and there was a stronger bias in the deep-learning-based CamCoder compared with the rule-based Edinburgh Geoparser. There was also strong regional variability in toponym resolution performance for highly

ambiguous toponyms, as revealed by the standard deviations of *MdnED* for annotated locations with the same name.

To probe the reason for the observed regional variability, we further evaluated geographic biases involved in geoparsing studies, ranging from spatial (diversity) misalignment of place coverage between (training and evaluation) corpora and the GeoNames gazetteer, the MAUP in toponym resolution, the preferences towards places with the largest population in place name disambiguation, to a biased perspective and usage of biased data in performance evaluation. Summing up, we rejected the claim that geoparsing is solved, and argued that it only appears so due to evaluation biases and that other geographic biases in geoparsing also need immediate attention.

6.2 Future Work

While we have discussed how recent work, such as spatial representation learning and spatially-explicit geoparsing performance evaluations, can help remove some of the geographic biases embedded in geoparsing, there is still a long way to go towards directions such as developing geoparsers that can succeed in performing accurately in place name disambiguation across geographic space. We hope to highlight aforementioned biases for future geoparsing research, and we call for debiasing work on geoparsing by following three potential paths.

First, more representative training and evaluation corpora is needed. When building such datasets, they should be evaluated by minimizing representation biases across geographic space. Also, the representativeness of toponyms in the context of a target toponym to be geoparsed should be assessed. An ablation study that masks either the target toponym feature or contextual toponym features showed degradation in the overall geoparsing performance [10], indicating a potential influence of the number, geographic

distribution or semantic characteristics of contextual toponyms on geoparsing. In addition, another critical data issue is the covariate shift caused by the difference between training and evaluation dataset distributions. This difference emerges because labeled articles in these datasets can be created differently during data sampling, thus preventing the scalability of geoparsing classifiers when applying them on evaluation datasets in the wild. To overcome this problem, a widely-adopted approach is to use pre-trained word embeddings (e.g., GloVe [64]). While they help improve the overall performance of geoparsing that can suffer from training data paucity, whether they correct the covariate shift caused by representation biases across geographic space remains to be studied.

Second, there are rich opportunities for developing more spatially-explicit geoparsers. Besides spatial representation learning on non-structured or semi-structured data, structured geographic information provided by knowledge graphs provide new possibilities in geoparsing [65, 66]. Combining graph representation learning [67] and spatial representation learning can be one of the many promising techniques that can extract additional geospatial semantics of places from interlinked data and embed them into geoparsing. However, a relevant ethical question is how to ensure algorithmic fairness towards underrepresented locations, such as those with no population information in gazetteers or with toponyms rarely mentioned in any corpus. As we also want to obtain a high geoparsing performance, we need to maintain a balance between geoparsing fairness and performance. Such maintenance requires both debiasing algorithms and quantitative measurements of geoparsing fairness.

Third, evaluation metrics used to evaluate geoparsing performance can incorporate more relevant spatial knowledge. Our evaluation has revealed geographic biases in toponym recognition, indicating that toponym recognition is also a geospatial task rather than merely a sub-task of NERC. This leads to a question of how to assess toponym recognition performance across geographic space by designing spatially-explicit evalua-

tion metrics. On the other hand, as direction is a first-class citizen of equal significance to distance [68], we may also need direction-based metrics in addition to distance-based metrics, particularly for toponym resolution performance evaluations. Such metrics can offer insights into anisotropy in geoparsing performance.

We also recommend further considerations about how much bias a task is able to bear when applying geoparsing in (geospatial) downstream tasks, e.g., vague cognitive region extraction [69] and geographic question answering [70]. Such considerations need to be taken into account in terms of dataset construction, algorithm design, and performance evaluations in geoparsing.

Bibliography

- [1] C. B. Jones and R. S. Purves, *Geographical information retrieval*, *International Journal of Geographical Information Science* **22** (2008), no. 3 219–228.
- [2] D. Nadeau and S. Sekine, *A survey of named entity recognition and classification*, *Linguisticae Investigationes* **30** (2007), no. 1 3–26.
- [3] S. Overell and S. Rüger, *Using co-occurrence models for placename disambiguation*, *International Journal of Geographical Information Science* **22** (2008), no. 3 265–287.
- [4] Y. Ju, B. Adams, K. Janowicz, Y. Hu, B. Yan, and G. McKenzie, *Things and strings: improving place name disambiguation from short texts by combining entity co-occurrence with topic modeling*, in *European Knowledge Acquisition Workshop*, pp. 353–367, Springer, 2016.
- [5] J. Wang, Y. Hu, and K. Joseph, *Neurotopr: A neuro-net toponym recognition model for extracting locations from social media messages*, *Transactions in GIS* **24** (2020), no. 3 719–735.
- [6] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, and C. Dyer, *Neural architectures for named entity recognition*, *arXiv preprint arXiv:1603.01360* (2016).
- [7] X. Hu, Z. Zhou, J. Kersten, M. Wiegmann, and F. Klan, *Gazpne2: A general and annotation-free place name extractor for microblogs fusing gazetteers and transformer models*, .
- [8] C. Zhou, C. Sun, Z. Liu, and F. Lau, *A c-lstm neural network for text classification*, *arXiv preprint arXiv:1511.08630* (2015).
- [9] M. Gritta, M. Pilehvar, and N. Collier, *Which melbourne? augmenting geocoding with maps*, .
- [10] S. Kulkarni, S. Jain, M. J. Hosseini, J. Baldrige, E. Ie, and L. Zhang, *Multi-level gazetteer-free geocoding*, in *Proceedings of Second International Combined Workshop on Spatial Language Understanding and Grounded Communication for Robotics*, pp. 79–88, 2021.

- [11] J. Fize, L. Moncla, and B. Martins, *Deep learning for toponym resolution: Geocoding based on pairs of toponyms*, *ISPRS International Journal of Geo-Information* **10** (2021), no. 12 818.
- [12] J. Wang and Y. Hu, *Are we there yet? evaluating state-of-the-art neural network based geoparsers using eupeg as a benchmarking platform*, in *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Geospatial Humanities*, pp. 1–6, 2019.
- [13] D. A. Griffith, *Spatial autocorrelation, A Primer* (Washington, DC, Association of American Geographers) (1987).
- [14] P. Legendre, *Spatial autocorrelation: trouble or new paradigm?*, *Ecology* **74** (1993), no. 6 1659–1673.
- [15] A. Getis, *A history of the concept of spatial autocorrelation: A geographer’s perspective*, *Geographical analysis* **40** (2008), no. 3 297–309.
- [16] R. R. Sokal and N. L. Oden, *Spatial autocorrelation in biology: 1. methodology*, *Biological journal of the Linnean Society* **10** (1978), no. 2 199–228.
- [17] W. D. Koenig, *Spatial autocorrelation of ecological phenomena*, *Trends in Ecology & Evolution* **14** (1999), no. 1 22–26.
- [18] A. Getis, *Reflections on spatial autocorrelation*, *Regional Science and Urban Economics* **37** (2007), no. 4 491–496.
- [19] Y. Hu and K. Janowicz, *An empirical study on the names of points of interest and their changes with geographic distance*, *arXiv preprint arXiv:1806.08040* (2018).
- [20] J. B. Dunning Jr, D. J. Stewart, B. J. Danielson, B. R. Noon, T. L. Root, R. H. Lamberson, and E. E. Stevens, *Spatially explicit population models: current forms and future uses*, *Ecological Applications* **5** (1995), no. 1 3–11.
- [21] E. G. Irwin and J. Geoghegan, *Theory, data, methods: developing spatially explicit economic models of land use change*, *Agriculture, Ecosystems & Environment* **85** (2001), no. 1-3 7–24.
- [22] D. L. DeAngelis and S. Yurek, *Spatially explicit modeling in ecology: a review*, *Ecosystems* **20** (2017), no. 2 284–300.
- [23] K. Janowicz, S. Gao, G. McKenzie, Y. Hu, and B. Bhaduri, *Geoai: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond*, 2020.
- [24] M. Goodchild, *Issues in spatially explicit modeling*, *Agent-based models of land-use and land-cover change* (2001) 13–17.

- [25] B. Yan, K. Janowicz, G. Mai, and R. Zhu, *xnet+ sc: Classifying places based on images by incorporating spatial contexts*, in *10th International Conference on Geographic Information Science (GIScience 2018)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [26] G. Mai, K. Janowicz, B. Yan, R. Zhu, L. Cai, and N. Lao, *Multi-scale representation learning for spatial feature distributions using grid cells*, in *International Conference on Learning Representations*, 2020.
- [27] S. Reddy and L. M. Dávalos, *Geographical sampling bias and its implications for conservation priorities in africa*, *Journal of Biogeography* **30** (2003), no. 11 1719–1727.
- [28] W. Yang, K. Ma, and H. Kreft, *Geographical sampling bias in a large distributional database and its effects on species richness–environment models*, *Journal of Biogeography* **40** (2013), no. 8 1415–1426.
- [29] M. M. Syfert, M. J. Smith, and D. A. Coomes, *The effects of sampling bias and model complexity on the predictive performance of maxent species distribution models*, *PloS one* **8** (2013), no. 2 e55158.
- [30] A. Basiri, M. Haklay, G. Foody, and P. Mooney, *Crowdsourced geospatial data quality: Challenges and future directions*, 2019.
- [31] K. Janowicz, Y. Hu, G. McKenzie, S. Gao, B. Regalia, G. Mai, R. Zhu, B. Adams, and K. Taylor, *Moon landing or safari? a study of systematic errors and their causes in geographic linked data*, in *The Annual International Conference on Geographic Information Science*, pp. 275–290, Springer, 2016.
- [32] J. Beck, M. Böller, A. Erhardt, and W. Schwanghart, *Spatial bias in the gbif database and its effect on modeling species’ geographic distributions*, *Ecological Informatics* **19** (2014) 10–15.
- [33] G. Zhang and A.-X. Zhu, *The representativeness and spatial bias of volunteered geographic information: a review*, *Annals of GIS* **24** (2018), no. 3 151–162.
- [34] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, *et. al.*, *Imagenet large scale visual recognition challenge*, *International journal of computer vision* **115** (2015), no. 3 211–252.
- [35] I. Krasin, T. Duerig, N. Alldrin, V. Ferrari, S. Abu-El-Haija, A. Kuznetsova, H. Rom, J. Uijlings, S. Popov, A. Veit, *et. al.*, *Openimages: A public dataset for large-scale multi-label and multi-class image classification*, *Dataset available from <https://github.com/openimages>* **2** (2017), no. 3 18.

- [36] S. Shankar, Y. Halpern, E. Breck, J. Atwood, J. Wilson, and D. Sculley, *No classification without representation: Assessing geodiversity issues in open data sets for the developing world*, *arXiv preprint arXiv:1711.08536* (2017).
- [37] A. S. Fotheringham and D. W. Wong, *The modifiable areal unit problem in multivariate statistical analysis*, *Environment and planning A* **23** (1991), no. 7 1025–1044.
- [38] W. Zeng, C. Lin, J. Lin, J. Jiang, J. Xia, C. Turkay, and W. Chen, *Revisiting the modifiable areal unit problem in deep traffic prediction with visual analytics*, *IEEE Transactions on Visualization and Computer Graphics* **27** (2020), no. 2 839–848.
- [39] E. Xu and H. Zhang, *Spatially-explicit sensitivity analysis for land suitability evaluation*, *Applied Geography* **45** (2013) 1–9.
- [40] D. N. Barton, S. Blumentrath, and G. Rusch, *Policyscape—a spatially explicit evaluation of voluntary conservation in a policy mix for biodiversity conservation in norway*, *Society & Natural Resources* **26** (2013), no. 10 1185–1201.
- [41] M. Forsius, H. Kujala, F. Minunno, M. Holmberg, N. Leikola, N. Mikkonen, I. Autio, V.-V. Paunu, T. Tanhuanpää, P. Hurskainen, *et. al.*, *Developing a spatially explicit modelling and evaluation framework for integrated carbon sequestration and biodiversity conservation: Application in southern finland*, *Science of the Total Environment* **775** (2021) 145847.
- [42] J. Cui, X. Kong, J. Chen, J. Sun, and Y. Zhu, *Spatially explicit evaluation and driving factor identification of land use conflict in yangtze river economic belt*, *Land* **10** (2021), no. 1 43.
- [43] J. Wang and Y. Hu, *Enhancing spatial and textual analysis with eupeg: An extensible and unified platform for evaluating geoparsers*, *Transactions in GIS* **23** (2019), no. 6 1393–1419.
- [44] M. Gritta, M. T. Pilehvar, and N. Collier, *A pragmatic guide to geoparsing evaluation*, *Language resources and evaluation* **54** (2020), no. 3 683–712.
- [45] C. Grover, *Lt-ttt2 example pipelines documentation*, *Edinburgh: Edinburgh Language Technology Group, July* **24** (2008).
- [46] G. DeLozier, J. Baldrige, and L. London, *Gazetteer-independent toponym resolution using geographic word profiles*, in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, 2015.
- [47] E. Kamaloo and D. Rafiei, *A coherent unsupervised model for toponym resolution*, in *Proceedings of the 2018 World Wide Web Conference*, pp. 1287–1296, 2018.

- [48] M. Gritta, M. T. Pilehvar, N. Limsopatham, and N. Collier, *What’s missing in geographical parsing?*, *Language Resources and Evaluation* **52** (2018), no. 2 603–623.
- [49] J. O. Wallgrün, M. Karimzadeh, A. M. MacEachren, and S. Pezanowski, *Geocorpora: building a corpus to test and train microblog geoparsers*, *International Journal of Geographical Information Science* **32** (2018), no. 1 1–29.
- [50] Y. Hu, K. Janowicz, and S. Prasad, *Improving wikipedia-based place name disambiguation in short texts using structured data from dbpedia*, in *Proceedings of the 8th workshop on geographic information retrieval*, pp. 1–8, 2014.
- [51] E. Laparra and S. Bethard, *A dataset and evaluation framework for complex geographical description parsing*, in *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 936–948, 2020.
- [52] C. Grover, R. Tobin, K. Byrne, M. Woollard, J. Reid, S. Dunn, and J. Ball, *Use of the edinburgh geoparser for georeferencing digitized historical collections*, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **368** (2010), no. 1925 3875–3889.
- [53] M. D. Lieberman, H. Samet, and J. Sankaranarayanan, *Geotagging with local lexicons to build indexes for textually-specified spatial data*, in *2010 IEEE 26th international conference on data engineering (ICDE 2010)*, pp. 201–212, IEEE, 2010.
- [54] M. Gritta, *Research data supporting” which melbourne? augmenting geocoding with maps”*, .
- [55] J. K. Ord and A. Getis, *Local spatial autocorrelation statistics: distributional issues and an application*, *Geographical analysis* **27** (1995), no. 4 286–306.
- [56] A. Getis and J. K. Ord, *The analysis of spatial association by use of distance statistics*, in *Perspectives on spatial data analysis*, pp. 127–145. Springer, 2010.
- [57] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, *A survey on bias and fairness in machine learning*, *ACM Computing Surveys (CSUR)* **54** (2021), no. 6 1–35.
- [58] G. Quattrone, L. Capra, and P. De Meo, *There’s no such thing as the perfect map: Quantifying bias in spatial crowd-sourcing datasets*, in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 1021–1032, 2015.
- [59] C. E. Shannon, *A mathematical theory of communication*, *The Bell system technical journal* **27** (1948), no. 3 379–423.

- [60] K. McDonough, L. Moncla, and M. Van de Camp, *Named entity recognition goes to old regime france: geographic text analysis for early modern french corpora*, *International Journal of Geographical Information Science* **33** (2019), no. 12 2498–2522.
- [61] K. Grossner and R. Mostern, *Linked places in world historical gazetteer*, in *Proceedings of the 5th ACM SIGSPATIAL International Workshop on Geospatial Humanities*, pp. 40–43, 2021.
- [62] S. Openshaw, *The modifiable areal unit problem*, *Quantitative geography: A British view* (1981) 60–69.
- [63] H. Suresh and J. V. Guttag, *A framework for understanding unintended consequences of machine learning*, .
- [64] J. Pennington, R. Socher, and C. D. Manning, *Glove: Global vectors for word representation*, in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- [65] M. Avvenuti, S. Cresci, L. Nizzoli, and M. Tesconi, *Gsp (geo-semantic-parsing): geoparsing and geotagging with machine learning on top of linked data*, in *European Semantic Web Conference*, pp. 17–32, Springer, 2018.
- [66] L. Nizzoli, M. Avvenuti, M. Tesconi, and S. Cresci, *Geo-semantic-parsing: Ai-powered geoparsing by traversing semantic knowledge graphs*, *Decision Support Systems* **136** (2020) 113346.
- [67] W. L. Hamilton, *Graph representation learning*, *Synthesis Lectures on Artificial Intelligence and Machine Learning* **14** (2020), no. 3 1–159.
- [68] R. Zhu, K. Janowicz, and G. Mai, *Making direction a first-class citizen of tobler’s first law of geography*, *Transactions in GIS* **23** (2019), no. 3 398–416.
- [69] S. Gao, K. Janowicz, D. R. Montello, Y. Hu, J.-A. Yang, G. McKenzie, Y. Ju, L. Gong, B. Adams, and B. Yan, *A data-synthesis-driven method for detecting and extracting vague cognitive regions*, *International Journal of Geographical Information Science* **31** (2017), no. 6 1245–1271.
- [70] G. Mai, K. Janowicz, R. Zhu, L. Cai, and N. Lao, *Geographic question answering: Challenges, uniqueness, classification, and future directions*, *AGILE: GIScience Series* **2** (2021) 1–21.