

# A Coherent Unsupervised Model for Toponym Resolution

Ehsan Kamaloo  
University of Alberta  
Edmonton, Alberta, Canada  
kamaloo@cs.ualberta.ca

Davood Rafiei  
University of Alberta  
Edmonton, Alberta, Canada  
drafie@cs.ualberta.ca

## ABSTRACT

Toponym Resolution, the task of assigning a location mention in a document to a geographic referent (i.e., latitude/longitude), plays a pivotal role in analyzing location-aware content. However, the ambiguities of natural language and a huge number of possible interpretations for toponyms constitute insurmountable hurdles for this task. In this paper, we study the problem of toponym resolution with no additional information other than a gazetteer and no training data. We demonstrate that a dearth of large enough annotated data makes supervised methods less capable of generalizing. Our proposed method estimates the geographic scope of documents and leverages the connections between nearby place names as evidence to resolve toponyms. We explore the interactions between multiple interpretations of mentions and the relationships between different toponyms in a document to build a model that finds the most coherent resolution. Our model is evaluated on three news corpora, two from the literature and one collected and annotated by us; then, we compare our methods to the state-of-the-art unsupervised and supervised techniques. We also examine three commercial products including Reuters OpenCalais, Yahoo! YQL Placemaker, and Google Cloud Natural Language API. The evaluation shows that our method outperforms the unsupervised technique as well as Reuters OpenCalais and Google Cloud Natural Language API on all three corpora; also, our method shows a performance close to that of the state-of-the-art supervised method and outperforms it when the test data has 40% or more toponyms that are not seen in the training data.

## CCS CONCEPTS

• **Information systems** → **Location based services; Information extraction; Geographic information systems; Probabilistic retrieval models;** • **Computing methodologies** → **Information extraction;**

## KEYWORDS

Toponym resolution, geolocation extraction, unsupervised disambiguation, context-bound hypotheses, spatial hierarchies

### ACM Reference Format:

Ehsan Kamaloo and Davood Rafiei. 2018. A Coherent Unsupervised Model for Toponym Resolution. In *WWW 2018: The 2018 Web Conference, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, Article 4, 10 pages. <https://doi.org/10.1145/3178876.3186027>

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

*WWW 2018, April 23–27, 2018, Lyon, France*

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5639-8/18/04.

<https://doi.org/10.1145/3178876.3186027>

## 1 INTRODUCTION

The size of the Web has been growing near-exponentially over the past decade with a vast number of websites emerging on a variety of subjects and large volumes of textual data being made available every day. In particular, a staggering amount of Web content (such as news articles, blog, forum posts, and tweets) that are added online on a minute by minute basis make frequent use of location names as points of reference. However, many place names have multiple interpretations and using them as references introduces ambiguity which in turn leads to uncertainty. Determining geographic interpretations for mentions of place names, known as *toponyms*, involves resolving multiple types of ambiguities. Toponym resolution is the task of disambiguating or resolving toponyms in natural language contexts to geographic locations (i.e., the corresponding lat/long values). One of the formidable challenges is therefore related to resolving the ambiguity of place names. For example, consider the word *Paris* in the following sentences:

- (1) “The November 2015 *Paris* attacks were the deadliest in the country since World War II.”<sup>1</sup>
- (2) “*Paris* was voted ‘the Prettiest Little Town in Canada’ by Harrowsmith Magazine.”<sup>2</sup>

The first sentence cites the tragic incidents in *Paris, France* while in the second sentence, the co-occurrence of *Canada* and *Paris* helps us identify *Paris*. This example illustrates that a toponym resolution method should probe for such clues in documents to reduce the inherent ambiguities of the natural language text. GeoNames<sup>3</sup>, the largest crowd-sourced location database, lists 97 interpretations for the place name *Paris*.

The problem of toponym disambiguation has been studied in the literature. Early works on geotagging documents rely on hand-crafted rules and heuristics (e.g., Web-a-Where [3]). Recent studies, however, are grounded on supervised and unsupervised models that do not warrant any manual rules [1, 8, 22, 27, 35]. Adaptive Context Features (or *Adaptive* in short), proposed by Lieberman and Samet [22], and TopoCluster, suggested by DeLozier et al. [8], are among the prominent methods that have been proposed in this area. Adaptive method casts toponym resolution as a classification problem, whereas TopoCluster leverages geographical measures to estimate geographical profiles for words.

In this paper, we propose an unsupervised model to tackle toponym resolution since supervised methods yield a poor performance due to the paucity of sufficient annotated data. Our methods rely merely on the document content and a gazetteer primarily because supplementary information about a Web document often is neither available nor reliable. Clearly, any additional data such

<sup>1</sup>[https://en.wikipedia.org/wiki/November\\_2015\\_Paris\\_attacks](https://en.wikipedia.org/wiki/November_2015_Paris_attacks)

<sup>2</sup><http://www.brant.ca/en/discover-brant/paris.asp>

<sup>3</sup><http://geonames.org/>

as the hosting site of the document and its location (if available) can further improve the performance.

Our toponym resolution model utilizes context-related features of documents. First, we develop a probabilistic model, called *Context-Bound Hypotheses (CBH)*, inspired by the work of Yu and Rafiei [37], to incorporate two context-related hypotheses into toponym resolution. Yu and Rafiei’s model aims at geotagging non-location entities and employs a primitive disambiguation technique to spatially resolve toponyms. We extend this model by integrating geographical information of locations into the hypotheses. These context-related premises capture some of the implicit relationships that hold between place names mentioned in the same document; thus, each toponym follows either the location of a frequent toponym or a nearby toponym. Then, we develop another model, called *Spatial-Hierarchy Sets (SHS)*, which discovers a minimal set of relationships (as discussed in Section 2) that can exist among toponyms. SHS maps the minimality problem to a conflict-free set cover problem wherein sets are constructed using containment and sibling relationships among toponyms. The final model, *Context-Hierarchy Fusion (CHF)*, merges CBH and SHS to exploit context features in extracting minimal relationships.

We conduct extensive experiments to evaluate our model. Our experiments are carried out on multiple datasets, one collected and annotated by us and two others well-known and used in the literature, covering a large range of news sources. We assess the performance of our model and compare it with the state-of-the-art supervised and unsupervised techniques as well as a few commercial geotagging products including Yahoo! YQL Placemaker<sup>4</sup>, Thomson Reuter’s OpenCalais<sup>5</sup>, and Google Cloud Natural Language API<sup>6</sup>. Moreover, we study the generalization problem of supervised methods by feeding unseen data to Adaptive classifier, showing that the classifier cannot keep up with our unsupervised model.

In summary, the key contributions of this work are as follows:

- We devise an unsupervised toponym resolution model that leverages context features of documents as well as spatial relationships of toponyms to produce a coherent resolution.
- We extensively evaluate our model on different datasets and in comparison with state-of-the-art methods.
- We demonstrate that our unsupervised model surpasses the state-of-the-art unsupervised technique, TopoCluster [8], and that it can handle unknown data better than supervised techniques.

The rest of this paper is structured as follows: Section 2 demonstrates a formal definition of the problem. The proposed unsupervised model is described in Section 3. Section 4 explains the evaluation criteria, gold corpora and experiments. Section 5 reviews the related work from the literature, and finally, in Section 6, we conclude with a summary of results.

## 2 PROBLEM DEFINITION

Given a document  $\mathcal{D}$  and a sequence of toponyms  $T = t_1, t_2, \dots, t_K$  mentioned in  $\mathcal{D}$  (e.g., extracted using a named-entity recognizer),

<sup>4</sup><https://developer.yahoo.com/yql/>

<sup>5</sup><http://www.opencalais.com/>

<sup>6</sup><https://cloud.google.com/natural-language/>

toponym resolution refers to grounding each toponym  $t_i$  to a geographic footprint  $\ell_i$  with a latitude and a longitude.

Geographic footprints or references are often derived from a gazetteer, a repository of georeferenced locations and their associated metadata such as type/class, population, spatial hierarchy, etc. Following previous works [22, 35], we select GeoNames as our gazetteer primarily because not only is it the largest public location database with sufficiently high accuracy [2], but it also stores the spatial hierarchy of locations<sup>7</sup>. Additionally, the bounding boxes of some locations can be retrieved from GeoNames.

Each toponym  $t_i$  in  $\mathcal{D}$  has a set of location interpretations  $L_i = \{l_{i,1}, l_{i,2}, \dots, l_{i,n_i}\}$ , derived from a gazetteer  $\mathcal{G}$ , where  $n_i$  indicates the number of interpretations for toponym  $t_i$ . Hence, toponym resolution can be seen as detecting a mapping from location mentions  $T$  to location interpretations. The resolution method yet cannot enumerate all possible combinations of interpretations. For instance, in a document that contains only 6 U.S. states: *Washington* ( $n_1=113$ ), *Florida* ( $n_2=228$ ), *California* ( $n_3=225$ ), *Colorado* ( $n_4=230$ ), *Arizona* ( $n_5=63$ ) and *Texas* ( $n_6=53$ ), the number of possible interpretations exceeds 4 billion. The past works in this area therefore incorporate heuristics to reduce the immense search space. For instance, picking the most populated interpretation is a simple heuristic that has been adopted in early works [19]. However, population alone cannot be effective for an off-the-shelf resolution system. We address this problem by looking into containment and sibling relationships among toponyms in a document.

## 3 THE UNSUPERVISED MODEL

The proposed method leverages a combination of context-related features of documents to address toponym resolution. These features are grounded on the characteristics of toponyms. It is well-accepted (e.g., SPIDER [35]) that toponyms mentioned in a document often show the following minimality properties:

- *one-sense-per-referent*: all of the occurrences of a toponym generally refer to a unique location within a single document;
- *spatial-minimality*: toponyms mentioned in a text tend to be in a spatial proximity of each other.

In this section, we develop context-bound hypotheses, inspired by the named entity geotagging method suggested by Yu and Rafiei [37]. Then, we describe spatial hierarchies built from containment and sibling relationships among location mentions in text. Lastly, we explain how these two methods coalesce into an unsupervised model to disambiguate toponyms.

### 3.1 Context-Bound Hypotheses

Yu and Rafiei [37] propose a probabilistic model to associate named entities to locations. The task of geotagging named entities is delineated as follows: given a named entity and a set of documents germane to it, a geotagger finds the geographic focus of the named entity. The model, introduced by Yu and Rafiei [37], incorporates two hypotheses: *geo-centre inheritance hypothesis* and *near-location hypothesis* and estimates the probabilities that these premises hold. The probabilistic model makes use of the known entities that are mentioned in the surrounding text to determine the geo-centre

<sup>7</sup>OpenStreetMap, another well-known crowd-sourced gazetteer, is ruled out since it does not contain spatial hierarchies [15].

of a named entity. Their geotagging task mainly focuses on non-location named entities and does only a simple location disambiguation on each toponym, independent of other toponyms in the same document. A question here is if their probabilistic model can be applied to toponym resolution. This is the question we study in our *Context-Bound Hypotheses (CBH)* model. In particular, to model the cohesion of toponyms to context, we integrate the hypotheses with geographical information of locations in order to spatially locate a place mention. Context-Bound assumptions allow us to reduce toponym resolution to a probabilistic model, which we are set to compute the estimations in this section.

---

**Algorithm 1:** Preliminary Toponym Disambiguation in CBH

---

**Input:** Document  $\mathcal{D}$  and sequence of toponyms  $T$   
**Output:** A preliminary mapping from  $T$  to location interpretations

```

1 resolution  $\leftarrow \emptyset$ 
2 for toponym  $t_i$  in  $T$  do
3    $\ell_i \leftarrow \text{NIL}$ 
4   for interpretation  $l_{i,j}$  in  $L_i$  do
5      $\mathcal{H}_{i,j} \leftarrow \text{RetrieveHierarchy}(l_{i,j})$ 
6     node  $\leftarrow \text{Lookup}(\text{parent}[l_{i,j}], \mathcal{H}_{i,j})$ 
7     score  $\leftarrow 0$ 
8     while node  $\neq \text{NIL}$  do
9       for  $m_h$  in Mentions(node) do
10        for  $m_l$  in Mentions( $l_{i,j}$ ) do
11          similarity  $\leftarrow \max(\text{similarity}, \frac{1}{\text{TD}(m_h, m_l)})$ 
12          score  $\leftarrow \text{score} + \text{similarity}$ 
13          node  $\leftarrow \text{parent}[\text{node}]$ 
14     if confidence[ $\ell_i$ ] < score then
15        $\ell_i \leftarrow (l_{i,j}, \text{score})$ 
16     else if confidence[ $\ell_i$ ] = score then
17       if population[ $\ell_i$ ] < population[ $l_{i,j}$ ] then
18          $\ell_i \leftarrow (l_{i,j}, \text{score})$ 
19 resolution  $\leftarrow \text{resolution} \cup (t_i, \ell_i)$ 

```

---

The space of possible interpretations (as shown with an example of 6 U.S. states) can be huge and enumerating all combinations may not be feasible. To be able to compute probabilities of the hypotheses, we perform a preliminary location disambiguation [37]. This procedure, which is shown in Algorithm 1, leverages a heuristic to resolve toponyms. Consider a location interpretation  $l_{i,j}$  of toponym  $t_i$ . The mentions of the ancestors in  $l_{i,j}$ 's spatial hierarchy (line 5; the hierarchies can be obtained from gazetteer  $\mathcal{G}$ ) can be used as clues to resolve toponym  $t_i$ . The closer an ancestor mention is, the more chance that particular interpretation has to get selected. For example, toponym *Edmonton* refers to 6 different locations. Provided that it co-occurred with either *Alberta* or *Canada*, we can pinpoint it (i.e., the city of *Edmonton* located in *Canada*). For each toponym  $t_i$ , the preliminary disambiguation measures a score for each interpretation  $l_{i,j}$  (lines 8-13) and picks the interpretation with maximum score (lines 14-15) and in case of

tie, the most populous interpretation is selected (lines 16-18). The score is acquired by finding the maximum similarity between  $l_{i,j}$  mentions and its ancestors' mentions; similarity here is the inverse of term distance (line 11), as used by Yu and Rafiei [37].

Preliminary disambiguation works poorly in cases where no mentions of locations in spatial hierarchy exist in the document. For instance, suppose we find toponyms *Toronto*, *London*, and *Kingston* in an article. Though, humans can recognize that these cities are presumably located in *Ontario, Canada*, preliminary resolution is unable to find any clues for disambiguation and as a result, assigns the toponyms to the interpretation with the highest population (i.e.,  $Toronto \mapsto Canada$ ,  $London \mapsto England$ , and  $Kingston \mapsto Jamaica$ ).

The result of the initial phase can be augmented by incorporating context-related features into the resolution process. Our CBH model proceeds to compute probabilities for the two hypotheses. The method operates at each administrative division separately since toponyms may lie in disparate division levels. Hence, the method begins the disambiguation process from the lowest division and furthers the process until all toponyms are resolved.

The geo-centre inheritance indicates that the location interpretation of a toponym can be drawn from the geographic scope of the document. The entities (i.e., people, locations, and organizations) used in an article, can ascertain a location interpretation to which the article is geographically relevant [4]. This location defines the geographic scope of the document.

Based on the inheritance hypothesis, the toponyms mentioned in a document are more likely to be part of or under the same administrative division as the geographic scope of the document. This makes sense due to the *spatial minimality* property. Therefore, we first estimate the geographic scope of the document via a probabilistic model. In particular, for toponym  $t_i$  at division  $d$ , the probability of  $l_{i,j}$  being the correct interpretation is

$$P_{\text{inh}}^{(d)}(l_{i,j}|\mathcal{D}, t_i) = \frac{\text{tf}(\text{anc}_d(l_{i,j}))}{\sum_{p=1}^{n_i} \text{tf}(\text{anc}_d(l_{i,p}))} \quad (1)$$

where  $\text{anc}_d$  returns the ancestor of an interpretation at division  $d$  and  $\text{tf}(w)$  computes the term frequency in the document. Each location interpretation here is extended to include its corresponding spatial hierarchy. For example, interpretations of toponym *Paris* are represented as

$$\{ [Paris \rightsquigarrow Ile-de-France \rightsquigarrow France], [Paris \rightsquigarrow LamarCounty \rightsquigarrow Texas \rightsquigarrow US], \dots \}$$

The second hypothesis, namely near-location hypothesis, relies upon the toponyms mentioned in the vicinity of a toponym. Toponyms nearby a toponym can be linked to one another primarily because of *object/container* and *comma group* relationships they possibly have [24]. According to this hypothesis, the closer toponym  $s$  to toponym  $t$ , the stronger evidence toponym  $s$  is to disambiguate toponym  $t$ . This is why, in this hypothesis, we compute the term distance between toponyms as a measure of similarity to estimate probabilities. In effect, for toponym  $t_i$  at division  $d$ , the probability of  $l_{i,j}$  being the correct interpretation is

$$P_{\text{near}}^{(d)}(l_{i,j}|\mathcal{D}, t_i) = \frac{\text{sim}(t_i, \text{anc}_d(l_{i,j}))}{\sum_{p=1}^{n_i} \text{sim}(t_i, \text{anc}_d(l_{i,p}))} \quad (2)$$

where  $\text{sim}(v_1, v_2)$  is the similarity function between terms  $v_1$  and  $v_2$  as demonstrated below:

$$\text{sim}(v_1, v_2) = \frac{1}{\min_{w_i \in M(v_i)} \{\text{TD}(w_1, w_2)\}} \quad (3)$$

where  $\text{TD}(w_1, w_2)$  is the distance between indices of  $w_1$  and  $w_2$  and  $M(v)$  is a set containing the mentions of term  $v$  in document  $\mathcal{D}$ .

Now, we combine  $P_{\text{inh}}^{(d)}$  and  $P_{\text{near}}^{(d)}$  to incorporate both premises into the model. The final context-bound model is regarded as a weighted linear function of the two probabilities:

$$P_{\text{CB}}^{(d)}(l_{i,j}|\mathcal{D}, t_i) = J^{(d)}(\mathcal{D}, t_i) \cdot P_{\text{near}}^{(d)}(l_{i,j}|\mathcal{D}, t_i) + (1 - J^{(d)}(\mathcal{D}, t_i)) \cdot P_{\text{inh}}^{(d)}(l_{i,j}|\mathcal{D}, t_i) \quad (4)$$

The coefficient  $J^{(d)}(\mathcal{D}, t_i)$  is obtained via Shannon Entropy of the vector induced by near-location probabilities for toponym  $t_i$  with respect to  $l_{i,j}$  for all values of  $j$ .

The resolution is undertaken through maximum likelihood estimation over the probability in Equation (4). The final computed probability can be considered as confidence score.

---

#### Algorithm 2: CBH Resolution

---

**Input:** Document  $\mathcal{D}$  and sequence of toponyms  $T$

**Output:** A mapping from  $T$  to location interpretations

```

1 resolution ← PreliminaryResol( $\mathcal{D}, T$ ) // Alg. 1
2 for  $k \leftarrow 1$  to  $\text{maxIterations}$  do
3   for division  $d$  in {County, State, Country} do
4     for toponym  $t_i$  in  $T$  do
5        $\ell_i \leftarrow \text{argmax}_j \{P_{\text{CB}}^{(d)}(l_{i,j}|\mathcal{D}, t_i)\}$ 
6       // refer to Eq. (4)
7       resolution ← resolution  $\cup (t_i, \ell_i)$ 

```

---

In summary, the CBH resolution method is illustrated in Algorithm 2. The approach starts with a preliminary resolution, followed by a hypotheses assessment to rectify results from the initial resolution. The hypotheses model computes the probabilities for each division separately to ensure the model can afford toponyms in all levels of dispersion. Once the modification process finished, the algorithm repeats for another iteration since altering the resolution of a toponym may affect other disambiguated toponyms. Our experiments show that CBH often takes two iterations to complete. However, in some cases, the modification step never terminates. Specifically, consider the following sentence, an excerpt from a news article:

“... London’s Heathrow, one of the world’s busiest travel hubs.”<sup>8</sup>

*London* and *Heathrow* are recognized as toponyms. Because no notion of ancestors in the spatial hierarchy can be found, the initial resolution favors the highest population interpretation (i.e.,  $\text{London} \mapsto \text{England}$  and  $\text{Heathrow} \mapsto \text{Florida, US}$ ). In the next step, the hypotheses model maps *London* to a place in United States because the other toponym is located in United States. Accordingly,

*Heathrow* is assigned to the airport in England. After the first iteration, the resolution is changed to  $\{\text{London} \mapsto \text{US}, \text{Heathrow} \mapsto \text{England}\}$ . Conversely, the second iteration would alter the results to  $\{\text{London} \mapsto \text{England}, \text{Heathrow} \mapsto \text{US}\}$ ; the algorithm is now trapped in an infinite loop. This is why, we introduce *maxIterations* parameter to eschew these circumstances. While CBH fails to successfully resolve toponyms in such cases, the approach, described in next section, can address this shortcoming.

### 3.2 Spatial-Hierarchy Sets

The spatial minimality property (noted by Leidner [19]) leads us to another resolution method called *Spatial-Hierarchy Sets (SHS)*. This method is grounded on containment and sibling relationships that are likely to exist among toponyms in a document. Consider a non-disjoint partitioning of the universe of locations (in a gazetteer) where locations with similar or related interpretations (e.g., those under the same administrative division or within a close proximity) form a partition. Since toponyms in a document tend to refer to geographically related locations, and those locations are more likely to be in the same partitions than different partitions, we want to find a small set of partitions that cover all toponyms; this can be modeled as a conflict-free covering problem. Conflict-free covering refers to the traditional set cover problem where each element must be covered by at most one set in the answer. The covering needs to be conflict-free due to *one-sense-per-referent* property. We formally define conflict-free covering as an instance of the conflict-free coloring of regions [16].

*Conflict-free Covering Problem.* Given a finite family of finite sets  $\mathcal{S}$  where each set  $S_i$  is associated with a non-negative weight  $w_i$  and a universal set  $\mathcal{U}$  containing all the elements from the sets, we seek to find a collection of sets, namely  $\mathcal{A}$ , with minimum weight such that their union becomes  $\mathcal{U}$  while each element is covered by at most one set in  $\mathcal{A}$ .

We formulate toponym resolution by conflict-free covering problem as the following:

- (1) Each parent with all its children form a set of related interpretations. Let  $\mathcal{S}$  denote the collection of all such sets that can be constructed. Each parent appears in a set with its children, hence the size of  $\mathcal{S}$  is the same as the number of parents with non-zero children. Algorithm 3 depicts the details of generating  $\mathcal{S}$ .
- (2) Recall that  $T$  denotes the set of toponyms in document  $\mathcal{D}$  (as defined in Section 2). We say a set in  $\mathcal{S}$  covers a toponym in  $T$ , if the set contains the surface text of the toponym. We want to select sets in  $\mathcal{S}$  that cover all toponyms in  $T$ . Our goal is to minimize the number of interpretations (spatial minimality) by selecting as few sets in  $\mathcal{S}$  as possible.
- (3) Let us form a color class for each toponym. The color class for a toponym includes all possible interpretations of the toponym. For example, *Texas* is a color class which includes all places that can resolve *Texas*. We want to avoid selecting multiple interpretations for the same toponym. That means, we do the selection in (2) with the constraint that no more than one color or interpretation can be selected for each toponym.

<sup>8</sup><http://money.cnn.com/2016/12/14/news/companies/british-airways-ba-strike-christmas>

---

**Algorithm 3:** Spatial-Hierarchy Set Generation

---

**Input:** Document  $\mathcal{D}$  and sequence of toponyms  $T$   
**Output:**  $\mathcal{S}$ , a collection of spatial hierarchy sets

```
1  $S \leftarrow \emptyset$ 
2  $P \leftarrow \emptyset$ 
3 for toponym  $t_i$  in  $T$  do
4   if  $name[t_i]$  in  $P$  then
5      $\text{skip } t_i$ 
6   for interpretation  $l_{i,j}$  in  $L_i$  do
7     // checks whether the set exists
8     if  $parent[l_{i,j}]$  in  $\mathcal{S}$  then
9        $\text{AddChild}((l_{i,j}, \text{true}), \mathcal{S}[parent[l_{i,j}]])$ 
10    else
11       $\mathcal{S} \leftarrow \mathcal{S} \cup \{(parent[l_{i,j}], \text{false}) \rightarrow (l_{i,j}, \text{true})\}$ 
12      // the new set is a tree rooted at
13      //  $parent[l_{i,j}]$ 
14      // the boolean values represent mentioned
15      // flags
16      if  $l_{i,j}$  in  $\mathcal{S}$  then
17         $\text{mentioned}[\mathcal{S}[l_{i,j}]] \leftarrow \text{true}$ 
18     $P \leftarrow P \cup name[l_{i,j}]$ 
```

---

In the special case where the color classes are empty (i.e., no constraint on colors), the problem becomes the classic set cover, which is NP-complete. This means that existing methods approximate the optimal solution. We leverage a greedy approach [34] to solve the problem. Although the greedy approach gives an approximate answer to the problem in general, our experiments reveal that such answer yield a competitive performance.

However, this model suffers from some deficiencies, even if an optimal solution is reached. A problem with this formulation is that we cannot have *Montreal, Quebec* and *Windsor, Ontario* in the same text (or they will not be resolved correctly) because *Windsor* is also a town in *Quebec*. These are cases where the hypotheses model, namely CBH, can better resolve. Furthermore, there may be circumstances that similar toponyms may appear in more than one sets and yet, we cannot favor one set to another. Suppose we have a document where only *Georgia* and *Turkey* are mentioned. Two sets,  $\{Georgia(\text{city}) \rightsquigarrow Texas(\text{state}), Turkey(\text{city}) \rightsquigarrow Texas(\text{state})\}$  and  $\{Georgia(\text{country}) \rightsquigarrow World, Turkey(\text{country}) \rightsquigarrow World\}$ , would emerge in  $\mathcal{S}$ . Without any additional information, such as document source, even humans cannot choose the correct interpretation. SHS selects the most populated set as a rule of thumb in these cases.

### 3.3 Context-Hierarchy Fusion

While the Spatial-Hierarchy Sets approach guarantees the minimality properties, it fails to select between identical structures (e.g., the *Georgia* and *Turkey* case) mostly because it does not delve into other context-related features of the document. On the other hand, the Context-Bound Hypotheses model benefits from term frequency and term distance features of the context. Notwithstanding

---

	TR-News	LGL	CLUST
News sources	36	85	352
Documents	118	588	13327
Annotated docs	118	588	1082
Annotated topos	1318	5088	11962
Topos with GeonameID	1274	4462	11567
Distinct topos	353	1087	2323
Median topos per doc	9	6	8
Topos not found in GeoNames	2.7%	3.2%	3.3%
Wikipedia-linked topos	94.3%	94.1%	94.2%

---

**Table 1: Corpora used in our experiments**

the situations like *Georgia* and *Turkey*, using other context sensitive information alleviates the disambiguation process in most cases. For example, toponyms *London, Aberdeen* and *Edinburgh* have interpretations located in *Canada* and SHS resolves them to the corresponding interpretations in *Canada* to preserve minimality. Even the presence of toponym *England* does not change the result because *Aberdeen* and *Edinburgh* located in *Scotland* and we still need to pick two sets to attain the correct resolution.

Consequently, merging SHS and CBH method allows us to take advantage of both methods at the same time. *Context-Hierarchy Fusion (CHF)* method chooses an interpretation from CBH only if the confidence score is higher than a threshold  $\tau$ . Otherwise, it resolves toponyms using SHS.

## 4 EXPERIMENTS

In this section, we conduct extensive experiments to evaluate our methods<sup>9</sup> and to assess their performance under different settings. The particular questions to be investigated are:

- (1) Given that CBH comprises different steps and components, how much does inheritance and near location hypothesis improve upon the preliminary location disambiguation?
- (2) How sensitive is Context-Hierarchy Fusion to the value of the threshold and if there are some sweet spots?
- (3) How accurate is the proposed method, compared to the state-of-the-art supervised and unsupervised methods as well as commercial systems?
- (4) How does the proposed method compare to the state-of-the-art supervised method in terms of the generality of the model on unseen data?
- (5) When is an unsupervised technique expected to surpass supervised methods?

For (3), we compare the performance of our method to that of the state-of-the-art methods as well as commercial systems (i.e., Yahoo! YQL Placemaker, OpenCalais and Google Cloud Natural Language API). The details of these proprietary products have not been made public. However, these systems can be accessed through public Web APIs at a relatively liberal rate limit, which enable us to automatically test their geotagging process on our datasets.

In our evaluation setting, we apply two methods for toponym recognition. First, we assume that the recognition phase is flawless,

---

<sup>9</sup>The source code and the annotated dataset is available at <https://github.com/ehsk/CHF-TopoResolver>

which is displayed as *Resol*. In this method, the annotated toponyms without latitude/longitude are fed to the underlying resolution method. These experiments are conducted to compare our methods to resolution methods such as TopoCluster [8]. Second, we employ Stanford NER [11] to tag locations, which is shown by *GeoTag*. We run *GeoTag* experiments to draw a comparison with systems performing both recognition and resolution including closed-source products and Adaptive [22].

#### 4.1 Datasets

In order to evaluate our toponym resolution methods, gold data corpora are required, in which all occurrences of geographic names and phrases have been manually annotated. In our experiments, we exploit three annotated datasets. Table 1 summarizes and compares the statistics of these datasets.

*TR-News*. We collected this dataset from various global and local News sources. We obtained news articles from several local news sources to include less dominant interpretations of ambiguous locations such as *Edmonton, England* and *Edmonton, Australia* rather than *Edmonton, Canada* or *Paris, Texas, US* in lieu of *Paris, France*. Additionally, a number of articles from global news sources such as BBC and Reuters have been selected to preserve the generality of the corpus. We manually annotated toponyms in the articles with the corresponding entries from GeoNames. The gold dataset consists of 118 articles.

*Local-Global Lexicon (LGL)*. This corpus was developed by Lieberman et al. [23]. It is collected from local news sources and mainly focuses on including ambiguous toponyms and this is why, it is suitable to test toponym resolution systems against geographically localized documents. The dataset is composed of 588 articles from 85 news sources.

*CLUST*. Lieberman and Samet [21] compiled this dataset from a variety of global and local news sources. CLUST is a large dataset containing 1082 annotated articles.

According to Table 1, the median number of toponyms per document in all datasets are close to each other, meaning that the corpora do not differ significantly with one another in terms of the number of toponyms per article.

In addition, the three datasets contain toponyms (roughly 3%) that cannot be found in gazetteer  $\mathcal{G}$ , while annotated and linked to an entry in the gazetteer. We observe that such toponyms fall into one of the following categories: uncommon abbreviations such as *Alta*, stands for *Alberta, Canada*, multi-word places such as *Montreal-Pierre Elliott Trudeau International Airport*, and transliterated place names (e.g., city of *Abbasiyeh, Egypt* written as *Abbasiya*).

The test corpora is also analyzed by the location type of their annotated toponyms, as done by Lieberman and Samet [22]. We compute the percentage of each location type for each dataset. As show in Figure 1, *LGL* dataset largely consists of small cities, which makes it a challenging test dataset since well-known locations are presumably to be resolved with high precision due to their frequent use in articles. In contrast, *TR-News* and *CLUST* datasets are roughly similar and include countries more than any location type.




Figure 1: Comparative analysis of the test datasets based on location type

This denotes that the articles appeared in *TR-News* and *CLUST* are extracted from sources that are aimed at a global audience. These sources usually provide more details for location mentions such as saying *Paris, US* instead of *Paris*. On the other hand, in *LGL*, because the articles are meant to be of use for local audience, the news publishers typically do not state additional information in this regard. Thus, geotagging approaches can be tested against these test corpora since they span a variety of news sources both globally and locally.

#### 4.2 Evaluation Metrics

Performance measures in our experiments are *Precision*, *Recall*,  $F_1$ -measure, and mean error distance ( $M$ ). However, to ascertain whether an interpretation is correctly predicted, we also investigate the error distance between the predicted coordinates and the actual coordinates, as used in numerous studies [5, 7, 8, 19, 22, 31, 35]. This distance enables us to fare various systems against each other since they may select latitude/longitude of locations from different gazetteers or knowledge bases<sup>10</sup>. We set the error distance to 10 miles, same as Adaptive method [22], whereas most researches tend to adopt a relaxed threshold (i.e., 161 kilometers) [5, 8, 31, 35].

For TopoCluster [8] and the commercial products, on the other hand, we employ a different criteria primarily because the error distance may not be accurate for large areas (even with higher error distance thresholds). Hence, in order to consider whether an interpretation is correctly projected to a coordinate, we check if the predicted interpretation resides in the bounding-box area of the ground truth; here the bounding-boxes of locations are extracted from GeoNames. We did not use this bounding-box grounded accuracy for other methods since most of them rely on the same gazetteer adopted in this work. Although using bound-boxes works in favor of TopoCluster and the proprietary products, the mean error distance metric fails to precisely mirror the accuracy for these methods since for a prediction deemed as correct based on bounding-boxes, the error distance can still be high.

Furthermore, in *Resol* experiments, we only calculate *Precision* because given a toponym, a resolution method is more likely to map it to an interpretation unless it does not exist in the gazetteer; thus, *Recall* would be approximately analogous to *Precision*. It is also

<sup>10</sup>Locations are represented with a single centroid and gazetteers may vary in picking the centroids.

	$P_{\text{Resol}}$	$P_{\text{GeoTag}}$	$R_{\text{GeoTag}}$	$F_{1-\text{GeoTag}}$
Preliminary	78.0	73.4	52.1	60.9
Inheritance	78.1	73.6	52.0	61.0
Near-location	79.0	73.9	52.3	61.2
CBH	79.2	74.9	53.0	62.1

**Table 2: Detailed analysis of Context-Bound Hypotheses (CBH) on *TR-News* dataset**

worth mentioning that the mean error distance is only reported in *Resol* experiments and not in *GeoTag* experiments, because the mean error distance cannot be measured for toponyms that are either not identified or falsely detected.

### 4.3 Analysis of Context-Bound Hypotheses

As discussed in Section 3.1, Context-Bound Hypotheses commences with a preliminary toponym disambiguation, followed by estimating two probabilities for inheritance and near-location hypotheses. In this section, we evaluate the preliminary phase and see whether the modification phase by Context-Bound hypotheses alleviates the resolution performance. Moreover, we study the role of the hypotheses in CBH by removing one of them at a time and measuring the performance. This experiment is conducted on the *TR-News* dataset in both *Resol* and *GeoTag* modes.

As shown in Table 2, taking both hypotheses into account complements the preliminary disambiguation, though the improvement does not seem considerable (slightly higher than 1% in  $F_1$ -measure) in both *GeoTag* and *Resol* experiments.

Additionally, the near-location hypothesis contributes to the improvement more than the inheritance hypothesis. This is largely because the inheritance hypothesis estimates probabilities using term frequency. In cases where two locations are mentioned as frequent as each other, term frequency does not seem accurate. For example, consider the toponym *Edmonton*, which can be located in either *Canada* or *Australia* in a document where *Australia* and *Canada* appear twice each. This results in the same score for both interpretations and a decision would be made by population size. Term distance, however, can help better in this case, denoting that the closer mention is more likely to be the correct interpretation. Nonetheless, we still need both hypotheses since the results are improved by putting near-location and inheritance together.

### 4.4 Fusion Threshold Study

In Context-Hierarchy Fusion (explained in Section 3.3), choosing an appropriate value for the threshold can be crucial in the resolution performance. In this experiment, we vary the threshold  $\tau$  to study its effect on performance. According to the results shown in Figure 2, we can identify a sweet spot when CHF achieves the best performance on all three datasets; this happens when  $\tau$  falls between 0.5 and 0.6; we set  $\tau$  to 0.55 in our experiments.

Also, we can see a mild spike in  $F_1$ -measure at  $\tau = 1$  in the *LGL* curve, which can be attributed to the localized content of the dataset. In particular, SHS (at  $\tau = 1$ , CHF is analogous to SHS) works better on *LGL* since locations in *LGL* are not mentioned frequently alongside their corresponding spatial hierarchy ancestors. As discussed in Section 3.1, CBH needs to spot the mentions of

these ancestors in documents (containment relationship) in order to generate a more accurate resolution, whereas SHS does not rely solely on containment relationships. It also takes sibling relationships into account, and as a result, merging SHS and CBH does not seem to be effective on *LGL*.

### 4.5 Resolution Accuracy

In this section, we measure the performance of our proposed methods and compare them with other resolution techniques. The methods presented in this paper are Context-Bound Hypotheses (CBH), Spatial-Hierarchy Sets (SHS) and Context-Hierarchy Fusion (CHF). We compare the results with two prominent systems: TopoCluster [8] (i.e., the state-of-the-art unsupervised model) and Adaptive [22] (i.e., the state-of-the-art supervised model). The source code of TopoCluster was available online, so we were able to test the method on our datasets. However, in order to test the Adaptive classifier, we implemented the supervised method<sup>11</sup>, albeit without two features, namely *dateline* and *locallex*; this was because for *locallex*, the authors used an annotated dataset containing the expected audience location of the news sources and also, *dateline* required a general location for each article which was not available for most articles in the test corpora. The modified version is named *CustomAdaptive* in our results. We follow the same parameter setting of the original Adaptive [22] and perform 10-fold cross validation to test CustomAdaptive.

Table 3 illustrates the evaluation results. CHF produces the best performance among our proposed methods on *CLUST* and *TR-News* and SHS beats the other proposed techniques on *LGL*. Among all listed methods, CustomAdaptive shows the highest performance. We also report recall, to make a comparison with the original Adaptive method [22].


While commercial products produce high precision, their recall is lower than our proposed methods in all cases except for Yahoo! YQL Placemaker. Placemaker yields the best results among the commercial products and achieves higher overall performance than our methods. On the other hand, OpenCalais is able to recognize toponyms as locative expressions. For instance, it identifies *the Kenyan captial* rather than just *Kenyan*. However, we observe that sometimes it fails to detect a full location phrase; for example, only *Toronto* in *Greater Toronto Area* is detected<sup>12</sup>. Further, Google Cloud Natural API offers an entity extraction service, which focuses highly on recognition of named entities<sup>13</sup>. The system links extracted entities to their corresponding Wikipedia articles and provides no additional information about geographic coordinates of location entities. Therefore, the geographical information of locations can only be derived from Wikipedia for this product. According to Table 1, nearly 94% of toponyms in each dataset have Wikipedia articles<sup>14</sup>, but not all Wikipedia articles contain spatial coordinates of locations, which is partly attributed to a poor recall in our experiments. Thus, we can see why entity linking approaches cannot be exploited for toponym resolution.

<sup>11</sup>Since we did not have access to the source code.

<sup>12</sup>We count these as correct resolutions unless they fall outside the bounding box of the annotated toponym.

<sup>13</sup>Google Cloud Natural API extracts locative expressions in any form in addition to proper names like *family home* and *suburb*.

<sup>14</sup>GeoNames keeps record of Wikipedia URLs for each location.




**Figure 2:**  $F_1$ -measure vs. threshold  $\tau$  for Context-Hierarchy Fusion method on *TR-News* dataset. At  $\tau = 0.55$ , CHF achieves the best  $F_1$ -measure on all three corpora.

We run *Resol* experiments to analyze TopoCluster [8] since it is a resolution method. DeLozier et al. stipulated that TopoCluster performs best when integrated with a gazetteer; this is why, the integrated version, called *TopoClusterGaz*, is adopted throughout this experiment. The results are presented in Table 3 ( $P_{\text{Resol}}$  and  $M_{\text{Resol}}$  columns). According to our results, CHF outperforms TopoCluster on all three datasets. Moreover, DeLozier et al. [8] set the error distance threshold for TopoCluster to 161 kilometers and achieved an accuracy of 71.4% on *LGL*<sup>15</sup>, whereas under the same setting, CHF reaches 71.2% on *LGL*, which is marginally lower than TopoCluster.

Besides accuracy, the mean error distance is also measured in our *Resol* experiments<sup>16</sup>. Among the unsupervised methods, CBH stands out with the lowest error. CHF is close to CBH with its error not exceeding 40km. This difference stems from SHS impacting CHF because when a toponym is projected to an incorrect location by SHS, the mapped location is more likely located in a country different than the ground truth.

#### 4.6 Unseen Data Analysis

Supervised techniques benefit from the knowledge gained in the training phase and if there is an overlap between the training data and the test data, then the prediction can be counted as overly optimistic. Domingos emphasizes that generalization is achieved by a separation of the training data and the test data [9]. This is why, we study the effect of the overlap between training and test datasets on  $F_1$ -measure. For this purpose, CustomAdaptive classifier was trained on *CLUST* dataset (the trend does not vary significantly if the classifier trained on *LGL*) and tested against *TR-News*. We define the overlap ratio measure as the number of toponyms per article in test data, which has also been appeared in the training data. We can channel overlap ratio through trimming off articles from test data and measure performance on the trimmed test data. Figure 3 plots  $F_1$ -measure against the overlap ratio. The unsupervised method surpasses the supervised method when the overlap ratio is less than 60% (when the overlap ratio is at 0.6, CHF still outperforms CustomAdaptive with a 1% margin). This observation confirms that the unsupervised technique, namely CHF, can handle



**Figure 3:**  $F_1$ -measure of CustomAdaptive trained on *CLUST* and CHF when overlap ratio varies. CHF yields a better performance than CustomAdaptive when overlap between training data and test data is lower than 60%.

unknown data better than the supervised method, namely Adaptive (CustomAdaptive implementation).

## 5 RELATED WORKS

Numerous studies have been conducted and much progress has been made on the task of disambiguating location mentions. The existing approaches in the literature may be grouped into (1) unsupervised and rule-based, (2) supervised, and (3) those based on some knowledge bases. However, a plethora of methods leverage a mixture of techniques. For example, DeLozier et al. [8] proposed an unsupervised toponym resolution method that leverages geographical kernels and spatially annotated Wikipedia articles. Also, Lieberman and Samet [22] presented a supervised technique that uses both geographical distance and additional knowledge like gazetteers and document source to disambiguate toponyms.

**Unsupervised and rule-based methods** In unsupervised resolution, various techniques have been studied. Map-based methods create a representation of all referents on a world map and apply techniques such as geographical centroid detection and outlier elimination to estimate the target of a toponym [19]. Moncla et al. [28] introduce a map-based technique where density-based clustering was carried out to detect outliers. Buscaldi [6] argues that map-based techniques face difficulties in grounding toponyms in a document when they are spatially far from each other. Rule-based and heuristic-based methods also have been adopted in the literature [3, 19]. For instance, the presence of “Canada” in text *London, Canada* may help disambiguate *London*. However, finding a set of rules to cover all cases in natural language text seems to be arduous.

**Approaches using knowledge bases** Wikipedia has been integrated as a knowledge base into more recent toponym disambiguation techniques [5, 8, 31, 35, 36]. Ardanuy and Sporleder [5] address toponym disambiguation in multilingual retrospective articles. They build a model to distill semantic features from Wikipedia information such as page title and article body. Speriosu and Baldrige [35] argue that non-spatial words impart useful information to disambiguate toponyms and they propose likelihood models that are obtained from Wikipedia. DeLozier et al. [8] propose TopoCluster, which does not rely on gazetteers to resolve toponyms, to address cases where location mentions are not found in gazetteers. They

<sup>15</sup> Among the datasets used in TopoCluster paper [8], *LGL* is the only dataset to which we have access

<sup>16</sup> Mean error distance for TopoCluster in *LGL* is derived from the original paper [8].



Method	LGL					CLUST					TR-News				
	P	R	F <sub>1</sub>	P <sub>Resol</sub>	M <sub>Resol</sub>	P	R	F <sub>1</sub>	P <sub>Resol</sub>	M <sub>Resol</sub>	P	R	F <sub>1</sub>	P <sub>Resol</sub>	M <sub>Resol</sub>
<b>Unsupervised</b>															
CBH	66.8	40.6	50.5	68.6	<b>760</b>	<b>80.6</b>	55.8	66.0	<b>81.5</b>	<b>709</b>	74.9	53.0	62.1	79.2	<b>869</b>
SHS	<b>69.7</b>	<b>43.3</b>	<b>53.4</b>	68.3	1372	72.8	51.6	60.4	71.1	1521	73.8	53.6	62.1	69.9	2305
CHF	68.5	43.1	52.9	<b>68.9</b>	818	<b>80.6</b>	<b>58.4</b>	<b>67.7</b>	81.0	788	<b>79.3</b>	<b>58.2</b>	<b>67.1</b>	<b>80.5</b>	942
TopoCluster [8]	-	-	-	59.7	1228	-	-	-	77.1	769	-	-	-	68.8	1422
<b>Supervised</b>															
Adaptive [22]	-	<b>58.7</b>	-	<b>94.2</b>	-	-	<b>61.8</b>	-	<b>96.0</b>	-	-	-	-	-	-
CustomAdaptive	<b>79.2</b>	48.5	<b>60.2</b>	88.3	<b>679</b>	<b>89.8</b>	57.9	<b>70.4</b>	93.4	<b>504</b>	<b>83.8</b>	<b>74.9</b>	<b>79.1</b>	<b>90.5</b>	<b>573</b>
<b>Commercial</b>															
Placemaker	73.5	<b>48.6</b>	<b>58.5</b>	-	-	87.4	<b>61.1</b>	<b>71.9</b>	-	-	80.8	<b>63.0</b>	<b>70.8</b>	-	-
OpenCalais	77.1	28.9	42.1	-	-	<b>87.5</b>	48.5	62.4	-	-	<b>81.3</b>	48.5	61.2	-	-
GoogleNL+Wiki	<b>80.5</b>	34.0	47.8	-	-	82.8	39.2	53.2	-	-	80.2	38.4	51.9	-	-

Table 3: Performance results of various methods in *GeoTag* and *Resol* experiments. The best results in each category are bolded.

construct a geographical language model to capture geographical senses of words using Wikipedia pages of locations. However, they note that adding gazetteer information to TopoCluster, namely *TopoClusterGaz*, yields a better performance. Less known toponyms are not expected to be found in Wikipedia; they can introduce challenges and hinder the performance of this method.

**Supervised methods** Many classification techniques have been proposed for geotagging purposes including Bayesian [1], random forests [22], RIPPER rule learner [14] and SVM [14, 27]. The features extracted for these classifiers can be grouped into context-free and context-sensitive features [22]. Context-free features typically include heuristics and information from external sources such as knowledge bases and gazetteers and may include, for example, population [22] and location type [14]. Context-sensitive features are obtained from documents where toponyms are mentioned. Melo and Martins [27] use normalized TF-IDF document vectors over curvilinear and quadrilateral regions on Earth’s surface. The adaptive method, proposed by Lieberman and Samet [22], casts geographical proximity and sibling relationship among interpretations in a context window as features. GeoWhiz [1] aggregate several likelihoods based on observations in training data. For instance, largely populated places are more likely estimated as their prominent interpretation. The suggested method by Santos et al. [31] consolidates information from Wikipedia pages of locations to compute several similarity and geographical features (context-free features) and performs a nearest neighbor search using locality-sensitive hashing (LSH) to resolve locations.

**Other more general related work** Entity disambiguation (also known as entity linking) [13, 17, 20, 25, 32] is related to toponym resolution. Linking named entities (i.e., people, organizations, and locations) to their corresponding real world entities in a knowledge base subsumes toponym disambiguation. Nonetheless, geographical features of location entities are neglected by these systems [35] and thus, geographically specialized methods for resolving toponyms are still needed to map locations to their corresponding geographic footprint.

Another line of research pertinent to this work is location disambiguation in social media. The related work in this area may

incorporate user profile data and social network information as well as natural language processing tools and gazetteers to tackle this task [18, 29]. Flatow et al. [12] propose a method that learns geo-referenced  $n$ -grams from training data to perform geotagging on social messages. Use of words that are endogenous to social media are considered as an inherent hurdle here. Moreover, social media content have deficient orthographic structure and lack context, which bring even more complexities to toponym resolution in social media [10, 26, 30].

## 6 CONCLUSIONS

In this paper, we study toponym resolution and propose two novel unsupervised models and a mixture model, namely CHF, to address the problem. We investigate the effectiveness of the proposed methods with other techniques. Our evaluations show that the Context-Hierarchy Fusion method outperforms TopoCluster, the state-of-the-art unsupervised method, in terms of precision. The performance of supervised techniques exceeds that of our proposed methods (as expected), nonetheless, we have shown that the state-of-the-art supervised classifier, called Adaptive, highly relies on the training data and Context-Hierarchy Fusion can handle unseen toponyms better.

The future work may investigate other mixture models and a better understanding of when one or both of supervised and unsupervised methods are expected to perform not so well. In addition, the correlations among the bounding-boxes of toponyms in an article can be studied to augment the resolution, considering the gazetteer are endowed with bounding-box of locations for this purpose [33]. Another direction is understanding the differences between short and long text as far as toponym resolution is concerned and the challenges each pose.

## ACKNOWLEDGMENTS

This research is supported by the Natural Sciences and Engineering Research Council of Canada. The authors are grateful to Michael Lieberman for sharing the *LGL* and *CLUST* datasets. The authors would also like to thank the anonymous referees for their valuable comments and helpful suggestions.

## REFERENCES

- [1] Marco D. Adelfio and Hanan Samet. 2013. GeoWhiz: Toponym Resolution Using Common Categories. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '13)*. ACM, New York, NY, USA, 532–535. <https://doi.org/10.1145/2525314.2525321>
- [2] Dirk Ahlers. 2013. Assessment of the accuracy of GeoNames gazetteer data. In *Proceedings of the 7th Workshop on Geographic Information Retrieval*. 74–81.
- [3] Einat Amitay, Nadav Har'El, Ron Sivan, and Aya Soffer. 2004. Web-a-where: Geotagging Web Content. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '04)*. ACM, New York, NY, USA, 273–280. <https://doi.org/10.1145/1008992.1009040>
- [4] G. Andogah, G. Bouma, and J. Nerbonne. 2012. Every Document Has a Geographical Scope. *Data Knowl. Eng.* 81–82 (Nov. 2012), 1–20. <https://doi.org/10.1016/j.datak.2012.07.002>
- [5] Mariona Coll Ardanuy and Caroline Sporleder. 2017. Toponym Disambiguation in Historical Documents Using Semantic and Geographic Features. In *Proceedings of the 2Nd International Conference on Digital Access to Textual Cultural Heritage (DATECH2017)*. ACM, New York, NY, USA, 175–180. <https://doi.org/10.1145/3078081.3078099>
- [6] Davide Buscaldi. 2011. Approaches to Disambiguating Toponyms. *SIGSPATIAL Special 3*, 2 (July 2011), 16–19. <https://doi.org/10.1145/2047296.2047300>
- [7] Zhiyuan Cheng, James Caverlee, and Kyumin Lee. 2010. You Are Where You Tweet: A Content-based Approach to Geo-locating Twitter Users. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM '10)*. ACM, New York, NY, USA, 759–768. <https://doi.org/10.1145/1871437.1871535>
- [8] Grant DeLozier, Jason Baldrige, and Loretta London. 2015. Gazetteer-independent Toponym Resolution Using Geographic Word Profiles. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, 2382–2388. <http://dl.acm.org/citation.cfm?id=2886521.2886652>
- [9] Pedro Domingos. 2012. A Few Useful Things to Know About Machine Learning. *Commun. ACM* 55, 10 (Oct. 2012), 78–87. <https://doi.org/10.1145/2347736.2347755>
- [10] Jacob Eisenstein. 2013. Phonological Factors in Social Media Writing. In *Proceedings of the Workshop on Language Analysis in Social Media*. Association for Computational Linguistics, Atlanta, Georgia, 11–19. <http://www.aclweb.org/anthology/W13-1102>
- [11] Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics (ACL '05)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 363–370. <https://doi.org/10.3115/1219840.1219885>
- [12] David Flatow, Mor Naaman, Ke Eddie Xie, Yana Volkovich, and Yaron Kanza. 2015. On the Accuracy of Hyper-local Geotagging of Social Media Content. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining (WSDM '15)*. ACM, New York, NY, USA, 127–136. <https://doi.org/10.1145/2684822.2685296>
- [13] Octavian-Eugen Ganea, Marina Ganea, Aurelien Lucchi, Carsten Eickhoff, and Thomas Hofmann. 2016. Probabilistic Bag-Of-Hyperlinks Model for Entity Linking. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 927–938. <https://doi.org/10.1145/2872427.2882988>
- [14] Eric Garbin and Inderjeet Mani. 2005. Disambiguating Toponyms in News. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Vancouver, British Columbia, Canada, 363–370. <http://www.aclweb.org/anthology/H/H05/H05-1046>
- [15] Judith Gelernter, Gautam Ganesh, Hamsini Krishnakumar, and Wei Zhang. 2013. Automatic gazetteer enrichment with user-geocoded data. In *Proceedings of the Second ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information*. ACM, 87–94.
- [16] Sarel Har-Peled and Shakhar Smorodinsky. 2003. On Conflict-free Coloring of Points and Simple Regions in the Plane. In *Proceedings of the Nineteenth Annual Symposium on Computational Geometry (SCG '03)*. ACM, New York, NY, USA, 114–123. <https://doi.org/10.1145/777792.777810>
- [17] Johannes Hoffart, Yasemin Altun, and Gerhard Weikum. 2014. Discovering Emerging Entities with Ambiguous Names. In *Proceedings of the 23rd International Conference on World Wide Web (WWW '14)*. ACM, New York, NY, USA, 385–396. <https://doi.org/10.1145/2566486.2568003>
- [18] Yohei Ikawa, Maja Vukovic, Jakob Rogstadius, and Akiko Murakami. 2013. Location-based Insights from the Social Web. In *Proceedings of the 22nd International Conference on World Wide Web (WWW '13 Companion)*. ACM, New York, NY, USA, 1013–1016. <https://doi.org/10.1145/2487788.2488107>
- [19] Jochen L. Leidner. 2007. Toponym Resolution in Text: Annotation, Evaluation and Applications of Spatial Grounding. *SIGIR Forum* 41, 2 (Dec. 2007), 124–126. <https://doi.org/10.1145/1328964.1328989>
- [20] Yang Li, Chi Wang, Fangqiu Han, Jiawei Han, Dan Roth, and Xifeng Yan. 2013. Mining Evidences for Named Entity Disambiguation. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '13)*. ACM, New York, NY, USA, 1070–1078. <https://doi.org/10.1145/2487575.2487681>
- [21] Michael D. Lieberman and Hanan Samet. 2011. Multifaceted Toponym Recognition for Streaming News. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 843–852. <https://doi.org/10.1145/2009916.2010029>
- [22] Michael D. Lieberman and Hanan Samet. 2012. Adaptive Context Features for Toponym Resolution in Streaming News. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '12)*. ACM, New York, NY, USA, 731–740. <https://doi.org/10.1145/2348283.2348381>
- [23] M. D. Lieberman, H. Samet, and J. Sankaranarayanan. 2010. Geotagging with local lexicons to build indexes for textually-specified spatial data. In *2010 IEEE 26th International Conference on Data Engineering (ICDE 2010)*. 201–212. <https://doi.org/10.1109/ICDE.2010.5447903>
- [24] Michael D. Lieberman, Hanan Samet, and Jagan Sankaranarayanan. 2010. Geotagging: Using Proximity, Sibling, and Prominence Clues to Understand Comma Groups. In *Proceedings of the 6th Workshop on Geographic Information Retrieval (GIR '10)*. ACM, New York, NY, USA, Article 6, 8 pages. <https://doi.org/10.1145/1722080.1722088>
- [25] Xiao Ling, Sameer Singh, and Daniel Weld. 2015. Design Challenges for Entity Linking. *Transactions of the Association for Computational Linguistics* 3 (2015), 315–328. <https://transacl.org/ojs/index.php/tacl/article/view/528>
- [26] Koji Matsuda, Akira Sasaki, Naoaki Okazaki, and Kentaro Inui. 2015. Annotating Geographical Entities on Microblog Text. In *Proceedings of The 9th Linguistic Annotation Workshop*. Association for Computational Linguistics, Denver, Colorado, USA, 85–94. <http://www.aclweb.org/anthology/W15-1609>
- [27] Fernando Melo and Bruno Martins. 2015. Geocoding Textual Documents Through the Usage of Hierarchical Classifiers. In *Proceedings of the 9th Workshop on Geographic Information Retrieval (GIR '15)*. ACM, New York, NY, USA, Article 7, 9 pages. <https://doi.org/10.1145/2837689.2837690>
- [28] Ludovic Moncla, Walter Renteria-Agualimpia, Javier Nogueras-Iso, and Mauro Gaio. 2014. Geocoding for Texts with Fine-grain Toponyms: An Experiment on a Geoparsed Hiking Descriptions Corpus. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '14)*. ACM, New York, NY, USA, 183–192. <https://doi.org/10.1145/2666310.2666386>
- [29] Sathappan Muthiah, Bert Huang, Jaime Arredondo, David Mares, Lise Getoor, Graham Katz, and Naren Ramakrishnan. 2015. Planned Protest Modeling in News and Social Media. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, 3920–3927. <http://dl.acm.org/citation.cfm?id=2888116.2888259>
- [30] Alan Ritter, Mausam, Oren Etzioni, and Sam Clark. 2012. Open Domain Event Extraction from Twitter. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '12)*. ACM, New York, NY, USA, 1104–1112. <https://doi.org/10.1145/2339530.2339704>
- [31] João Santos, Ivo Anastácio, and Bruno Martins. 2015. Using machine learning methods for disambiguating place references in textual documents. *GeoJournal* 80, 3 (01 Jun 2015), 375–392. <https://doi.org/10.1007/s10708-014-9553-y>
- [32] W. Shen, J. Wang, and J. Han. 2015. Entity Linking with a Knowledge Base: Issues, Techniques, and Solutions. *IEEE Transactions on Knowledge and Data Engineering* 27, 2 (2015), 443–460. <https://doi.org/10.1109/TKDE.2014.2327028>
- [33] Sanket Kumar Singh and Davood Rafiei. 2018. Strategies for geographical scoping and improving a gazetteer. In *WWW 2018: The 2018 Web Conference*. ACM, New York, NY, USA. <https://doi.org/10.1145/3178876.3186078>
- [34] Petr Slavik. 1996. A Tight Analysis of the Greedy Algorithm for Set Cover. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing (STOC '96)*. ACM, New York, NY, USA, 435–441. <https://doi.org/10.1145/237814.237991>
- [35] Michael Speriosu and Jason Baldrige. 2013. Text-Driven Toponym Resolution using Indirect Supervision. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Sofia, Bulgaria, 1466–1476. <http://www.aclweb.org/anthology/P13-1144>
- [36] Andreas Spitz, Johanna Geiß, and Michael Gertz. 2016. So Far Away and Yet So Close: Augmenting Toponym Disambiguation and Similarity with Text-based Networks. In *Proceedings of the Third International ACM SIGMOD Workshop on Managing and Mining Enriched Geo-Spatial Data (GeoRich '16)*. ACM, New York, NY, USA, Article 2, 6 pages. <https://doi.org/10.1145/2948649.2948651>
- [37] Jiangwei Yu and Davood Rafiei. 2016. Geotagging Named Entities in News and Online Documents. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management (CIKM '16)*. ACM, New York, NY, USA, 1321–1330. <https://doi.org/10.1145/2983323.2983795>