

Knowledge Enabled Approach to Predict the Location of Twitter Users

Revathy Krishnamurthy*, Pavan Kapanipathi*, Amit P. Sheth, and
Krishnaprasad Thirunarayan

Kno.e.sis Center, Wright State University, USA,
{revathy,pavan,amit,tkprasad}@knoesis.org

*Joint First Authors

Abstract. Knowledge bases have been used to improve performance in applications ranging from web search and event detection to entity recognition and disambiguation. More recently, knowledge bases have been used to analyze social data. A key challenge in social data analysis has been the identification of the geographic location of online users in a social network such as Twitter. Existing approaches to predict the location of users, based on their tweets, rely solely on social media features or probabilistic language models. These approaches are supervised and require large training dataset of geo-tagged tweets to build their models. As most Twitter users are reluctant to publish their location, the collection of geo-tagged tweets is a time intensive process. To address this issue, we present an alternative, knowledge-based approach to predict a Twitter user’s location at the city level. Our approach utilizes Wikipedia as a source of knowledge base by exploiting its hyperlink structure. Our experiments, on a publicly available dataset demonstrate an improvement of 3% in the accuracy of prediction and a 16% reduction in the average error distance, over the state of the art supervised techniques.

Keywords: Wikipedia, Twitter, Location Prediction, Semantics

1 Introduction

Location of Twitter users is a prominent attribute for many applications such as emergency management and disaster response [13], trend prediction [1], and event detection [19]. Twitter users can choose to publish their location information by way of (1) geo-tagging their tweets, or (2) specifying it in the location field of their Twitter profile. However, recent studies have shown that less than 4% of tweets are geo-tagged [11, 15]. Also, while many users choose to leave the location field of their profile empty or enter invalid information, others specify location at different granularity such as city, state, and country. Thus, most of the information entered in this field cannot be reverse geocoded to a city. For instance, Cheng et al. [4] found that, in their dataset comprising of 1 million Twitter users, only 26% of the users shared their location at the city level.

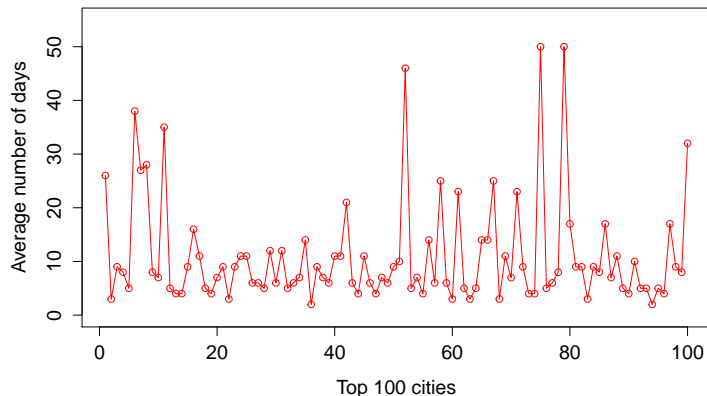


Fig. 1. Estimates of average number of days to collect geo-tagged tweets for top 100 cities in the training dataset of [4]

Existing approaches to predict the location of Twitter users, based on their tweets, use supervised learning techniques [3, 4]. They are built on the hypothesis that the geographic location of users influences the content of their tweets. These approaches are data-driven and require large training dataset of geo-tagged tweets to build statistical models that predict a user’s location. Cheng et al. [4] created a training dataset comprising of 4,124,960 geo-tagged tweets from 130,689 users in continental United States. The collection of this dataset was time intensive and done over a period of 5 months from September 2009 to January 2010. However, in the recent times we have seen a rapid growth of Twitter. Hence, we examined the effort required to create a similar data set in the present day. We selected the top 100 cities with the maximum count of tweets in the dataset of [4] and collected geo-tagged tweets from these cities over a period of 5 days. Based on the tweets collected in this duration, Figure 1 shows the average number of days required to collect tweets comparable in volume to [4]. We see that for some cities it would take up to 50 days for creating a high quality training data set. This makes it a time intensive process; consequently, making the approach challenging to adapt to newer cities. We address this weakness by proposing a knowledge based solution.

Knowledge bases have been used to either propose alternatives to learning approaches [9, 10] or in combination with learning approaches to improve their performance [7, 20]. This work falls in the former category. Our approach can be organized into three steps: (1) First, the *Creation of a Location Specific Knowledge base*, which exploits the hyperlink structure of Wikipedia to build a knowledge base of location specific entities. Additionally, we weight each entity by its ability to discriminate between locations. (2) Second, *User Profile Generation*,

which creates a semantic profile of a Twitter user whose location is to be determined. The user profile consists of Wikipedia entities found in their tweets and are weighted to reflect their importance to the user. (3) Finally, we use the overlap between the entities in the tweets of a user and the location specific knowledge base to rank the top k locations of the user in the *Knowledge base enabled Location Prediction* step. Concretely, we make the following contributions:

- We propose a novel knowledge based approach to predict the location of a Twitter user at the city level.
- We introduce the concept of *Local Entities* which are entities that can discriminate between geographic locations.
- We evaluate our approach using a benchmark dataset published by Cheng et al. [4] and show that our approach performs better than the state of the art approaches.

The rest of the paper is organized as follows. Section 2 describes the creation of a location specific knowledge base. Section 3 describes our approach to predict the location of a user using the location specific knowledge base. Section 4 describes the evaluation and results of our approach. Section 5, we explain the related work on location prediction. Finally, Section 6 concludes with suggestions for future work.

2 Creation of Location Specific Knowledge base

To create a location specific knowledge base, we (1) identify the local entities of a city, and (2) compute their localness measure with respect to the city.

2.1 Local Entities

Previous research that address the problem of location prediction of Twitter users, have established that the content of a user’s posts reflects his/her location. Cheng et al. [4] introduced the idea of *Local Words* which are words that convey a strong sense of location. For example, they found that the word *rockets* is local to Houston whereas words such as *world* and *peace* are more generic and do not exhibit an association to any particular location. Based on the same intuition, we introduce the concept of *local entities*. Local entities are Wikipedia entities that can distinguish between locations.

We leverage Wikipedia to identify the local entities for each city. While there are many knowledge bases, such as Yago¹, DMOZ², and Geo Names³, we choose Wikipedia because (1) it is comprehensive, (2) it contains dedicated pages for cities, and (3) it has a hyperlink structure that can be exploited for our purposes. A Wikipedia page comprises of links to other Wikipedia pages. These links are

¹ <http://www.mpi-inf.mpg.de/yago-naga/yago>

² <http://www.dmoz.org>

³ <http://download.geonames.org/export/dump/>

referred to as *internal links*⁴ of the Wikipedia page and are semantically related to the page (or a portion of it) [17]. Consequently, we consider the entities represented by internal links in the Wikipedia page of a city as local entities of that city. For example, the Wikipedia page of *San Francisco* contains a link to the Wikipedia page of *Golden Gate Bridge*. Thus, we consider *Golden Gate Bridge* as a local entity with respect to *San Francisco*.

2.2 Localness Measures

All the local entities of a city are not equally local with respect to city. For example, consider *San Francisco Giants* and *Major League Baseball* that are local entities of the city *San Francisco*⁵. While the *San Francisco Giants* are a baseball team based out of San Francisco, *Major League Baseball* is a professional baseball organization in North America. Intuitively, the entity *San Francisco Giants* has a higher potential than *Major League Baseball* to distinguish *San Francisco* from other cities in United States. Therefore, we introduce the concept of localness measure for each local entity such that the localness score reflects the distinguishing ability of the local entity with respect to a city. We experiment with four measures to determine the localness of an entity. These measures can be classified into three categories: (1) association based measure, (2) graph based measure, and (3) semantic overlap based measures.

Association Based Measure In information theory, pointwise mutual information is a standard measure of association. It is used to determine association between terms based on the probability of their co-occurrence. The intuitive basis for using an association measure to establish localness of an entity is that, higher the co-occurrence of a local entity with the city in Wikipedia pages, higher is the localness of the entity with respect to the city. In order to determine the association between a local entity and a city, we utilize the whole Wikipedia corpus. We define the PMI of a city and its local entity as:

$$PMI(le, c) = \log_2 \frac{P(le, c)}{P(c)P(le)} \quad (1)$$

where c is the city and le is a local entity of the city.

We compute the joint probability of occurrence, $P(c, le)$ as the fraction of the Wikipedia pages that contain both the city and the entity. Additionally, the individual probabilities of the city $P(c)$ and the local entity $P(le)$ are computed as the fraction of the Wikipedia pages that contain the city alone and the local entity alone respectively.

Graph Based Measure The Wikipedia hyperlink structure can also be represented as a directed graph whose vertices are the Wikipedia pages. An edge in

⁴ <http://en.wikipedia.org/wiki/Help:Link#Wikilinks>

⁵ http://en.wikipedia.org/wiki/San_Francisco

this graph represents a link from the Wikipedia page of the source node to the Wikipedia page of the target node. Since the hyperlink structure of Wikipedia allows us to represent a city and its local entities as a graph, we use a graph theoretic measure to compute the localness of the local entities.

To construct the graph of local entities for a city, we prune the Wikipedia hyperlink graph by selecting only those edges that connect the local entities of the city. For instance, *San Francisco Giants* and *Major League Baseball* are nodes in the graph of local entities of San Francisco. A directed edge from *San Francisco Giants* to *Major League Baseball* represents the link from the Wikipedia page of *San Francisco Giants* to the Wikipedia page of *Major League Baseball*.

Betweenness centrality has been used extensively to find influential nodes in a network. Our hypothesis is that, the relative importance of a node in the graph of local entities, reflects the localness of the local entity with respect to the city. It is defined as follows:

$$C_B(le, c) = \sum_{le_i \neq le \neq le_j} \frac{\sigma_{le_i le_j}(le)}{\sigma_{le_i le_j}} \quad (2)$$

where c is a city, le, le_i, le_j are local entities of c . $\sigma_{le_i le_j}$ represents the total number of shortest paths from le_i to le_j and $\sigma_{le_i le_j}(le)$ is the number of shortest paths from le_i to le_j through le . We normalize the measure by dividing C_B by $(n-1)(n-2)$ where n is the number of nodes in the directed graph.

Semantic Overlap Measure Halaschek et al. [8] measure the relatedness between concepts with the hypothesis that related concepts are connected to similar entities. Similarly, we measure the localness of a n entity with respect to a city as the overlap between the internal links of the entity and the internal links of the city. To compute this semantic overlap, we use the following set based measures: (1) Jaccard Index, and (2) Tversky Index .

Jaccard Index is a symmetric measure of overlap between two sets and is normalized for their sizes. Jaccard Index for a city c and its local entity le is defined as follows:

$$jaccard(le, c) = \frac{|IL(c) \cap IL(le)|}{|IL(c) \cup IL(le)|} \quad (3)$$

where $IL(c)$ and $IL(le)$ are the internal links found in the Wikipedia page of city c and local entity le respectively.

Tversky Index is an asymmetric measure of overlap of two sets [21]. While the Jaccard Index determines the overlap between a city and a local entity, a local entity generally represents a part of the city. For example, consider the local entity *Boston Red Sox*⁶ of the city *Boston*. Its internal links may not symmetrically overlap with that of *Boston* because internal links of *Boston* are from different categories such as *Climate*, *Geography* and *History*. Hence, we

⁶ Boston Red Sox is the baseball team of Boston

adapt Tversky Index to measure unidirectional overlap of the local entity le to the city c as follows:

$$ti(le, c) = \frac{|IL(c) \cap IL(le)|}{|IL(c) \cap IL(le)| + \alpha|IL(c) - IL(le)| + \beta|IL(le) - IL(c)|} \quad (4)$$

where we choose $\alpha = 0$ and $\beta = 1$ to penalize the local entity, for every internal link in its page not found in the Wikipedia page of the city.

3 Knowledge enabled Location Prediction

We create a semantic profile of the user whose location is to be predicted. It comprises of Wikipedia entities mentioned in their tweets. From this profile, entities that are local entities of a city are used to rank the top-k locations of a user.

3.1 User Profile Generation

Our approach is based exclusively on the contents of a user’s tweets. The *User Profile Generator* can be explained in two steps: (1) Entity Recognition from user’s tweets; (2) Entity Scoring to measure the extent of the usage of the entity by the Twitter user.

Entity Recognition Entity recognition is the process of recognizing information like people, organization, location, and numeric expressions⁷ from twitter messages. In this paper, our main focus is on the location prediction of Twitter users. Hence, we use the APIs available for Entity Recognition. We opted for Zemanta because: (1) It has been shown to be superior to others as evaluated against other entity recognition and linking services, by Derczynski et al. [5]; (2) Zemanta’s web service⁸ also links entities from the tweets to their Wikipedia pages. This allows an easy mapping between the Zemanta annotations and our knowledge base extracted from Wikipedia; (3) It provides co-reference resolution for the entities; (4) Zemanta increased their API rate limit to 10,000 requests per day for research purposes.⁹

Entity Weighting We weight each entity with the frequency of the occurrence of the entity in the user’s tweets. Frequency of mentions of an entity indicates the significance of the entity to the user.

⁷ More details on entity recognition can be found in [16]

⁸ <http://developer.zemanta.com/docs/suggest/>

⁹ We thank Zemanta for their support

3.2 Location Prediction

To predict the location of a user, we compute a score for each city whose local entities are found in the profile of the user, defined as follows:

$$locScore(u, c) = \sum_{e \in LE_{cu}} locl(e, c) \times s_e \quad (5)$$

where LE_{cu} is the set of local entities of c found in the profile of user u , $locl(e, c)$ is the localness measure of the entity e with respect to the city c and s_e is the weight of the local entity in the user profile. The location of the user is determined by ranking the cities in the descending order of $locScore(c, u)$.

4 Evaluation

First, we compare our approach with the four localness measures explained in Section 2.2. Then, we use the best performing measure to evaluate against the state of the art tweet based location prediction algorithms.

4.1 Dataset

For a fair comparison of our approach against the existing approaches, we use the dataset published by Cheng et al. [4]. The dataset contains 5119 users, from the continental United States, with approximately 1000 tweets of each user. A user location is listed in the form of latitude and longitude coordinates which is generally more reliable than the profile information. Spammers and bots are filtered to ensure a clean dataset.

To create the location specific knowledge base, we consider all the cities of United States with population greater than 5000, as published in the census estimates of 2012. We extract the local entities of a city from its Wikipedia page available in the XML dump of Wikipedia¹⁰. Finally, our location specific knowledge base comprises of 4,661 cities with 500,714 local entities.

4.2 Evaluation-Metrics

We adopt the following four evaluation measures used by the existing location prediction approaches [4]:

- Accuracy (ACC): The percentage of users identified within 100 miles of their actual location.
- Average Error Distance (AED): The average of the error distance across all users. *Error distance* is the distance between the actual location of the user and the estimated location by our algorithm.

¹⁰ http://en.wikipedia.org/wiki/Wikipedia:Database_download

- *Accuracy@k* (ACC@k): The percentage of users whose actual locations are within the *top-k* predicted locations of the user, with an error distance of 100 miles.
- *Average Error Distance@k* (AED@k): The Average of error distance, between the closest predicted location@*top-k* to the actual location, across all the users in the dataset.

4.3 Baseline

We implement a baseline system which considers all the entities of a city to be equally local to the city. To predict the location of a user, we compute the score for each city by aggregating the count of local entities of the city found in the user’s tweets and selecting the city with the maximum score. In other words, the localness measure of each entity in Equation 5 is 1.

4.4 Results

Method	ACC	AvgErrDist (in Miles)	ACC@2	ACC@3	ACC@5
Baseline	25.21	632.56	38.01	42.78	47.95
PMI	38.48	599.408	49.85	56.06	64.15
BC	47.91	478.14	57.39	62.18	66.98
JC	53.21	433.62	67.41	73.56	78.84
TI	54.48	429.00	68.72	74.68	79.99

Table 1. Location Prediction using Local Entities

Location Prediction Using Local Entities: Table 1 reports the Accuracy and the Average Error Distance for location prediction using the (1) Baseline, (2) Pointwise Mutual Information (PMI), (3) Betweenness Centrality (BC), (4) Semantic Overlap Measures - Jaccard Index (JC), and (5) Semantic Overlap Measures - Tversky Index (TI) .

We see that Tversky Index is the best performing localness measure with approximately 55% accuracy and 429 miles of AED. The accuracy is doubled compared to the baseline approach. However, compared to Jaccard Index, there is only a slight improvement in accuracy from 53.21% to 54.48% and decrease in AED from 433 to 429 miles.

The top *k* cities for a user are determined by ordering the aggregate score for each city (defined in Equation 5). As shown in Figure 2, the accuracy of our approach increases with *k*. At *k* = 5, using Tversky Index as the localness measure, we are able to predict the location of approximately 80% of the users. Similarly, as shown in Figure 3, the average error distance decreases as *k* increases.

Figure 4 shows the accuracy of prediction within increasing radius (in miles). As seen in the graph, we can predict approximately 46% of the users within 30 miles of the actual location of the user.

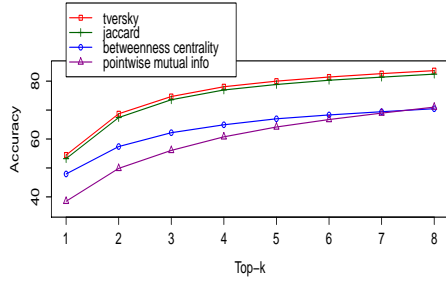


Fig. 2. Top-K accuracy

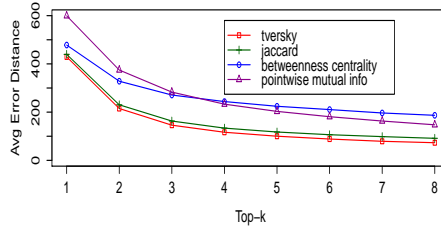


Fig. 3. Top-K Average Error Distance

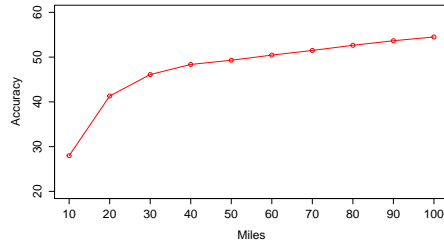


Fig. 4. Accuracy of prediction at increasing miles of radius

Table 2 shows examples of local entities extracted from the tweets of users. These examples illustrate that local entities of various types such as sports team, landmarks, organizations and famous people are used to predict the location of a user.

Performance of Localness Measures: Table 1 and Figures 2 and 3 have compared the results of our approach using the localness measures described in Section 2.2. In this section, we discuss our findings and analysis on why some localness measures performed better than others in the location prediction task.

Pointwise mutual information measure is sensitive to low frequency data [?]. This led to high absolute PMI scores for the local entities of a city like *Glen Rock, New Jersey* as compared to that of *San Francisco* due to the low occurrence of

City	Entities
New York City, NY	New York City, Brooklyn, Harlem, Queens, New York Knicks, The Bronx, Manhattan, National Football League, American Broadcasting Company, Train station, Metro-North Railroad, Rapping, Times Square, Fox Broadcasting Company, Broadway theatre, New York Yankees, Staten Island, Brooklyn Nets, Hudson River, Macy’s Thanksgiving Day Parade
Houston, TX	Houston; Houston Texans; NASA; Houston Astros; Interstate 45; Houston Chronicle; Greater Houston; Harris County, Texas; Galveston, Texas; Downtown Houston; Houston Rockets; Texas
Seattle, WA	Seattle; Seattle Seahawks; Seattle metropolitan area; Kobe; Microsoft; Downtown Seattle; Light rail; Alki Point
Nashville, TN	Nashville, Tennessee; Belmont University; Frist Center for the Visual Arts; Southeastern Conference; Centennial Park (Nashville); Gaylord Opryland Resort & Convention Center
Pittsburgh, PA	Pittsburgh; Midwestern United States; PNC Park; Station Square; Squirrel Hill (Pittsburgh); Giant Eagle; Fort Pitt Tunnel; Pittsburgh Steelers; Luke Ravenstahl; University of Pittsburgh;

Table 2. Examples of Local Entities found in tweets

the former as compared to the latter in the Wikipedia corpus. Nevertheless, the prediction results using PMI show a significant improvement over the baseline.

Betweenness centrality, as a localness measure, performs better than PMI. Although, betweenness centrality addresses the sensitivity to low frequency, it weights certain generic entities higher than more specific entities. For example, in our knowledge base, *United States* is a local entity with respect to the city *San Francisco*. We found that there are multiple shortest paths through *United States* in the graph of local entities of *San Francisco*, thus increasing its importance. However, the entity *United States* is fairly generic and cannot discriminate between the cities in our knowledge base.

The semantic overlap measures overcome the disadvantages of both betweenness centrality and PMI. The primary distinction between the two semantic overlap measures is that Jaccard Index is symmetric while Tversky Index is asymmetric. Jaccard Index is biased against local entities that have less internal links. For example, consider the two entities *Eureka Valley*, *San Francisco* and *California*. Both are local entities of the city *San Francisco*. Intuitively, we would expect *Eureka Valley*, *San Francisco* (a residential neighbourhood in San Francisco) to be more local than *California* with respect to the city *San Francisco* but with Jaccard Index the result is opposite. This problem motivated the use of an asymmetric measure. Using the Tversky Index, the localness measure of an entity is the highest when all its internal links are subsumed by those in the Wikipedia page of the city. Furthermore, the local entity is penalized for only the

internal links in its page not present in the city. Therefore, in the above example it is able to assign a higher degree of localness to *Eureka Valley, San Francisco* than *California* with respect to the city *San Francisco*. This approach to weighting the local entities performs better than Jaccard’s index with improved accuracy and lower average error distance.

Comparison with Existing Approaches: For the location prediction task based on user’s tweets, the state of the art approaches require a training dataset of geo tagged tweets. Their models are trained using a dataset of 4.1 million tweets collected over 5 months between 2009 and 2010. From Figure 1, we can see that the collection of geo-tagged tweets for the *top-100* cities (ranked based on the number of geo-tagged tweets from the cities used to train the models [3, 4]) in 2015 can take upto 50 days¹¹. On the other hand, our approach requires a pre-processing step of creating an index of the wikipedia links that can be easily downloaded from dbpedia¹². Once the index is created it is a matter of minutes to create the location specific knowledge base in comparison to tens of days required for the supervised techniques.

We compare the existing approaches against our approach with the best performing localness measure, i.e. Tversky Index (see Table 3). For a fair comparison in the results, we have evaluated our approach on the same test dataset as Cheng et al. [4] and Chang et al. [3]. As reported in Table 3, our approach performs 3.48% better in terms of the accuracy compared to the best performing approach. Furthermore, the average error distance is reduced by approximately 16%, i.e around 80 miles.

Method	ACC	AvgErrDist (in Miles)
Cheng et al. 2010 [4]	51.00	535.564
Chang et al. 2012 [3]	49.9	509.3
Our approach with TI	54.48	429.00

Table 3. Location prediction results compared to existing approaches

Impact of Local Entities on Prediction Accuracy: Twitter users have varying frequency of mentioning local entities of their city in tweets. From the test dataset of 5119 users (with approximately 1000 tweets per user), Figure 5 shows the percentage of users against the frequency of *distinct local entity* mentions. While 40% of users have mentioned less than 5 distinct local entities, 25% of the users have more than 10 distinct local entities in their tweets. The impact in determining the location of users based on this varying frequency of

¹¹ This experiment was performed keeping in mind the extensive growth of Twitter from 2009 to 2015.

¹² <http://wiki.dbpedia.org/Downloads38>

local entity mentions is shown in Figure 6. The accuracy of prediction increases with increase in the number of distinct local entities in the tweets of a user. The accuracy of prediction is 66% for users who mention more than 10 local entities in their tweets. As a part of our future work, we intend to include this analysis to contribute to a confidence score of our predictions.

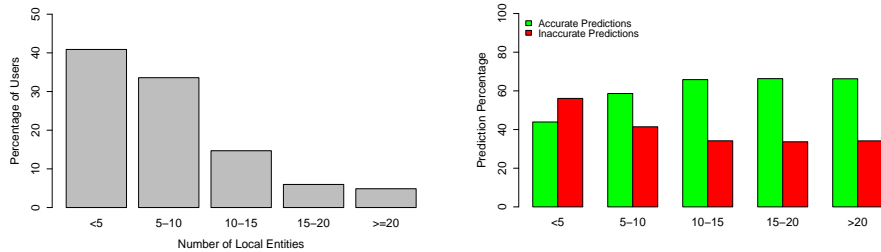


Fig. 5. Percentage of users with the count **Fig. 6.** Predictions corresponding to the of distinct local entities from the predicted count of distinct local entities in users' city tweets (Localness measure: Tversky Index)

Geographic Distribution of Predictions: The count of local entities for the cities in our knowledge base ranges from 11 (for *Island Lake, Illinois*) to 1095 (for *Chicago*). This reflects the information available in Wikipedia about the city. The first thought is that, the availability of information on Wikipedia can impact the performance of our approach. In order to analyze this issue, we performed experiments to check if any bias exists towards specific cities based on the number of local entities in the knowledge base. Hence, we plotted the distribution of our accurate predictions on a map of United States (Figure 7) against the distribution of all test users in the dataset as shown in Figure 8. We can see that despite the variation in the amount of information available for each city, our algorithm was able to predict locations of users from 356 distinct cities. The knowledge base for these accurately predicted cities ranged between of 40 to 1095 local entities. To conclude, we see that the accurate prediction (within 100 miles) is not restricted to few cities.

5 RELATED WORK

Geo-locating twitter users has gained a lot of traction due to its potential applications. Existing approaches to solve this problem can be grouped in to classes: (1) content based location prediction, and (2) network based location prediction.

Content-based location prediction approaches are grounded on the premise that the online content of a user is influenced by their geographical location. It

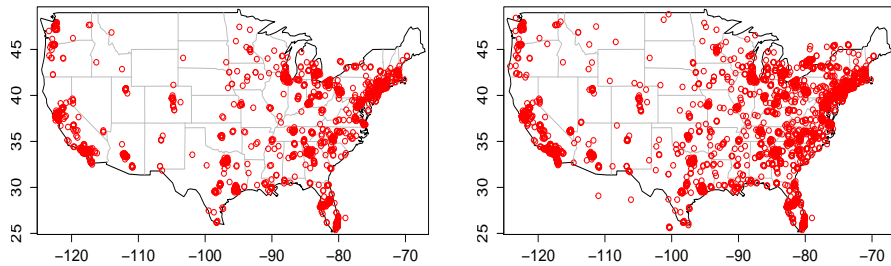


Fig. 7. Distribution of users predicted within 100 miles of their location **Fig. 8.** Distribution of all users in the dataset

relies on a significantly large training dataset to build a statistical model that identifies words with a local geographic scope. Cheng et al. [4] proposed a probabilistic framework for estimating a Twitter user’s city-level location based on the content of approximately 1000 tweets of each user. They formulated the task of identifying local words as a decision problem. They used the model of spatial variation proposed by [2] to train a decision tree classifier using a hand-curated list of 19,178 words. Their approach on a test dataset of 5119 users, could locate 51% of the users within 100 miles with an average error distance of 535 miles. The disadvantage of this approach was the assumption that a “term” is spatially significant to or characteristic of only one location/city. This challenge was addressed by Chang et al. [3] by modeling the variations as a Gaussian mixture model. While this approach still required a training dataset of geo-tagged tweets, it did not need a labeled set of seed words. Their tests on the same dataset showed an accuracy (within 100 miles) of 49.9% with 509.3 miles of average error distance. Eisenstein et al. [6] proposed cascading topic models to identify lexical variation across geographic locations. Using the regional distribution of words, determined from these models, they predicted the locations of twitter users. Their dataset comprised of users from United States. Their accuracy at the region and state level was 58% and 27% respectively. Kinsella et al. [12] addressed two problems, namely, (1) predicting the location of an individual tweet and (2) predicting the location of a user. They created language models for each location at different granularity levels of country, state, city and zipcode, by estimating a distribution of terms associated with the location.

Network based solutions are grounded in the assumption that the locations of the people in a user’s network and their online interaction with the user can be used to predict his/her location. McGee et al. [14] used the interaction between users in a network to train a Decision Tree to distinguish between pairs of users likely to live close by. They reported an accuracy of 64% (within 25 miles). Rout et al. [18] formulated this task as a classification task and trained an SVM classifier with features based on the information of users’ followers-followees who have their location information available. They tested their approach on a random sample of 1000 users and reported 50.08% accuracy at the city level.

However, a network based approach can only be used to determine the location of users who have other users in their network whose location is already known.

6 CONCLUSION AND FUTURE WORK

In this paper, we presented a novel knowledge based approach that uses Wikipedia to predict the location of Twitter users. We introduced the concept of *Local Entities* for each city and demonstrated the results of different measures to compute the localness of the entities with respect to a city. Without any training dataset, our approach performs better than the state of the art content based approaches. Furthermore, our approach can expand the knowledge base to include other cities which is remarkably less laborious than creating and modeling a training dataset.

In future, we will explore the use of semantic types of the Wikipedia entities to improve the accuracy of the location prediction and decrease the average error distance. We also plan to augment our knowledge base with location information from other knowledge bases such as Geo Names and Wikitravel. Additionally, we will examine how to adapt our approach to predict the location of a user at a finer granularity level like the neighborhoods in a city.

References

1. Harshavardhan Achrekar, Avinash Gandhe, Ross Lazarus, Ssu-Hsin Yu, and Benyuan Liu. Predicting Flu Trends using Twitter Data. In *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*, pages 702–707. IEEE, 2011.
2. Lars Backstrom, Jon Kleinberg, Ravi Kumar, and Jasmine Novak. Spatial Variation in Search Engine Queries. In *Proceedings of the 17th international conference on World Wide Web*, pages 357–366. ACM, 2008.
3. Hau-wen Chang, Dongwon Lee, Mohammed Eltaher, and Jeongkyu Lee. @ phillies Tweeting from Philly? Predicting Twitter User Locations with Spatial Word Usage. In *ASONAM 2012*, 2012.
4. Zhiyuan Cheng, James Caverlee, and Kyumin Lee. 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*, pages 759–768. ACM, 2010.
5. Leon Derczynski, Diana Maynard, Niraj Aswani, and Kalina Bontcheva. Microblog-Genre Noise and Impact on Semantic Annotation Accuracy. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*, pages 21–30. ACM, 2013.
6. Jacob Eisenstein, Brendan O’Connor, Noah A Smith, and Eric P Xing. A Latent Variable Model for Geographic Lexical Variation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1277–1287. Association for Computational Linguistics, 2010.
7. Evgeniy Gabrilovich and Shaul Markovitch. Overcoming the Brittleness Bottleneck using Wikipedia: Enhancing Text Categorization with Encyclopedic Knowledge. In *AAAI*, volume 6, pages 1301–1306, 2006.

8. Chris Halaschek, Boanerges Aleman-Meza, I Budak Arpinar, and Amit P Sheth. Discovering and Ranking Semantic Associations over a Large rdf Metabase. In *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*, pages 1317–1320. VLDB Endowment, 2004.
9. Xiaohua Hu, Xiaodan Zhang, Caimei Lu, Eun K Park, and Xiaohua Zhou. Exploiting Wikipedia as External Knowledge for Document Clustering. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 389–396. ACM, 2009.
10. Pavan Kapanipathi, Prateek Jain, Chitra Venkataramani, and Amit Sheth. User Interests Identification on Twitter Using a Hierarchical Knowledge Base. In *The Semantic Web: Trends and Challenges*, volume 8465 of *Lecture Notes in Computer Science*, pages 99–113. Springer International Publishing, 2014.
11. Saurabh Khanwalkar, Marc Seldin, Amit Srivastava, Anoop Kumar, and Sean Colbath. Content-based Geo-Location Detection for Placing Tweets Pertaining to Trending News on Map. In *The Fourth International Workshop on Mining Ubiquitous and Social Environments*, page 37, 2013.
12. Sheila Kinsella, Vanessa Murdock, and Neil O’Hare. I’m Eating a Sandwich in Glasgow: Modeling Locations With Tweets. In *Proceedings of the 3rd international workshop on Search and mining user-generated contents*, pages 61–68. ACM, 2011.
13. Kirill Kireyev, Leysia Palen, and K Anderson. Applications of Topics Models to Analysis of Disaster-related Twitter Data. In *NIPS Workshop on Applications for Topic Models: Text and Beyond*, volume 1, 2009.
14. Jeffrey McGee, James Caverlee, and Zhiyuan Cheng. Location Prediction in Social Media Based on Tie Strength. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, pages 459–468. ACM, 2013.
15. Fred Morstatter, Jürgen Pfeffer, Huan Liu, and Kathleen M Carley. Is the Sample Good Enough? Comparing Data from Twitters Streaming API with Twitters Firehose. *Proceedings of ICWSM*, 2013.
16. David Nadeau and Satoshi Sekine. A Survey of Named Entity Recognition and Classification. *Linguisticae Investigationes*, 30(1):3–26, 2007.
17. Andrea Giovanni Nuzzolese, Aldo Gangemi, Valentina Presutti, and Paolo Ciancarini. Encyclopedic Knowledge Patterns from Wikipedia Links. In *The Semantic Web-ISWC 2011*, pages 520–536. Springer, 2011.
18. Dominic Rout, Kalina Bontcheva, Daniel Preotjiuc-Pietro, and Trevor Cohn. Where’s @wally?: A Classification Approach to Geolocating Users Based on their Social Ties. In *HT*, 2013.
19. Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors. In *Proceedings of the 19th international conference on World wide web*, pages 851–860. ACM, 2010.
20. Yangqiu Song, Haixun Wang, Zhongyuan Wang, Hongsong Li, and Weizhu Chen. Short Text Conceptualization using a Probabilistic Knowledgebase. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume Three*, pages 2330–2336. AAAI Press, 2011.
21. Amos Tversky. Features of Similarity. *Psychological review*, 84(4):327, 1977.