

# Geotagging One Hundred Million Twitter Accounts with Total Variation Minimization

Ryan Compton, David Jurgens, David Allen  
Information and System Sciences Laboratory  
HRL Laboratories  
3011 Malibu Canyon Rd, Malibu, CA 90265.  
rfcompton@hrl.com

**Abstract**—Geographically annotated social media is extremely valuable for modern information retrieval. However, when researchers can only access publicly-visible data, one quickly finds that social media users rarely publish location information. In this work, we provide a method which can geolocate the overwhelming majority of active Twitter users, independent of their location sharing preferences, using only publicly-visible Twitter data.

Our method infers an unknown user’s location by examining their friend’s locations. We frame the geotagging problem as an optimization over a social network with a total variation-based objective and provide a scalable and distributed algorithm for its solution. Furthermore, we show how a robust estimate of the geographic dispersion of each user’s ego network can be used as a per-user accuracy measure which is effective at removing outlying errors.

Leave-many-out evaluation shows that our method is able to infer location for 101,846,236 Twitter users at a median error of 6.38 km, allowing us to geotag over 80% of public tweets.

**Keywords**-Social and Information Networks; Data mining; Optimization

## I. INTRODUCTION

The ability to geospatially index a large volume of Twitter data is valuable for several emerging research directions. Indeed, geographic social media analytics have proven useful for understanding regional flu trends [21], linguistic patterns [18], election forecasting [28], social unrest [6], and disaster response [16]. These approaches, however, depend on the physical locations of Twitter users which are only sparsely available in public data.

Interestingly, recent work from the computational social sciences community has established that online social ties are often formed over short geographic distances [27] [19] [8]. Because of this, it is possible to approximate the location of a Twitter user by examining publicly-known locations of their online friends [11] [31]. Using social network analysis to solve geolocation problems relies only on public Twitter metadata, which, somewhat counterintuitively, provides several advantages over content-based approaches. Network-based geotagging sidesteps difficulties with foreign natural language processing, ignores noisy Twitter text, and makes it possible to demonstrate results at previously unreachable scales. The largest and most accurate geolocation results

currently published have utilized social network analysis [11] [1].

In this work, we solve network-based geotagging problems from an optimization viewpoint. Unlike existing methods which have independently developed node-wise heuristics, we show that a globally-defined and highly-studied convex optimization can be solved for location inference. To be precise, we infer user location by solving

$$\min_{\mathbf{f}} |\nabla \mathbf{f}| \text{ subject to } f_i = l_i \text{ for } i \in L \quad (1)$$

where  $\mathbf{f} = (f_1, \dots, f_n)$  encodes a location estimate for each user,  $L$  denotes the set of users who opt to make their locations,  $l_i$ , public, the total variation on the Twitter social network is defined by

$$|\nabla \mathbf{f}| = \sum_{ij} w_{ij} d(f_i, f_j) \quad (2)$$

where  $d(\cdot, \cdot)$  measures geodesic distance via Vincenty’s formulae, and the edge weights,  $w_{ij}$ , are equal to the minimal number of reciprocated @mentions between users  $i$  and  $j$ .

In other words, **we seek a network such that the weighted sum over all geographic distances between connected users is as small as possible**. This sum, defined in (2) and known in the literature as total variation, has demonstrated superior performance as an optimization heuristic for several information inference problems across a wide variety of fields.

The minimization in (1) leads to a network where connected communities of users are placed into the same city. Motivation for this heuristic relies on communicative locality in Twitter which has been demonstrated in previous research [27] [11] [31] and is further justified by our experiments in section III-B. Evidence of communicative locality in other online social networks can be found in [19] [8] and [1].

The social media geotagging problem is relatively new; total variation minimization, however, is not. Originally introduced to the computational sciences as a heuristic for image denoising [23], literature on total variation minimization is now truly vast — thousands of papers spanning several decades. While most of this research has been confined to imaging processing, recent work has shown that

total variation is valuable to several newer applications, such as crime modelling [26], graphics [20], and “community detection” [10]. Very recently, it has been shown that use of total variation for general transductive learning tasks outperforms state-of-the-art methods on standard benchmark datasets [3] [2]. The present work is the first to demonstrate a theory link between the social media geotagging problem and total variation minimization.

There are real and immediate practical applications for our results. From a business perspective, the importance of Twitter geotagging has recently warranted several commercial offerings<sup>12</sup>. At least one product<sup>3</sup> is based on alignments between self-reported profile locations and the GeoNames<sup>4</sup> database. Our experiments in section IV-A demonstrate that, even if all nonempty profile locations could be mapped to unambiguous locations, the coverage of profile-based geotagging could reach at most 63% of tweet volume. When nonsensical or ambiguous profile locations are ignored, our experiments show that the coverage of profile-based geotagging drops below 20%.

A salient feature of our approach, however, is that we are unable to estimate within-city motion for individual Twitter users. Our goal is high-volume static location inference with city-level accuracy. This may be tolerable for several applications as research has indicated that users typically remain within a small radius; cf. section III-F and [4]. Since it is still unclear if better-than-city-level accuracy is possible at scale, we restrict our training sets to users who tweet primarily within the same radius, using their median location as “home”

Using social networks to infer location makes sense only if a user’s friends are primarily located within the same geographic region. We show that the geographic dispersion of each user’s ego network<sup>5</sup>,

$$\tilde{\nabla} f_i = \text{median}_j (w_{ij} d(f_i, f_j)) \quad (3)$$

agrees well with geotag error and therefore provides us with a per-user confidence measure. Overall error can thus be controlled by limiting  $\max \tilde{\nabla} f_i$ . Our experiments show that coverage remains high even when restrictions on  $\max \tilde{\nabla} f_i$  are tight.

Solutions to the geotagging problem are most interesting when demonstrated at scale. We exhibit our method on a network of 110,893,747 Twitter users with 1,034,362,407 connections between them and infer locations for 101,846,236 users. This is made possible by the optimization algorithm we employ to solve (1), parallel

coordinate descent [22], which we have implemented in a distributed manner using the Apache Spark cluster computing framework [32].

To be clear, the contributions of this paper are as follows:

- We show that the social network geotagging problem can be solved via a globally-defined convex optimization
- We develop a novel per-user accuracy estimate
- We demonstrate our results at scale; producing the largest database of Twitter user locations currently known in the literature

## II. BACKGROUND AND RELATED WORK

Location inference in social media has caught the attention of several researchers from both academia and industry. Two major classes of solutions to the geotagging problem are now prevalent: language-based, and network-based.

### A. Language-based Geotagging

Work by [5] and [7] provided methods based on identifying and searching for location-specific terms in Twitter text. More recent work by [13] (extended in [14]) builds on these approaches with ensemble classifiers that account for additional features such as time zone and volume of tweets per hour. To the best of our knowledge, [14] showcases the current state-of-the-art for content-based geotagging: 9,551 test users at 68% city-accuracy. Our goal is similar to that of [14] in that both works focus on static location inference with city-resolution accuracy, however, we demonstrate several orders of magnitude greater coverage as well as higher accuracy. Specifically, our leave-many-out validation tests indicate that we can correctly predict the city a user resides in with vs 89.7% city-accurate vs. the 68% reported by [14].

Geotagging via natural language processing requires that users from different geographic regions tweet in different dialects, and that these differences are great enough to make accurate location inference possible. Evidence against this possibility appears in [11], where the author examined a large Twitter dataset and found minimal agreement between language models and proximity. Additionally, language-based geotagging methods often rely on sophisticated language-specific natural language processing and are thus difficult to extend worldwide.

### B. Network-based Geotagging

Large-scale language-agnostic geotagging is possible by inferring a user’s location with the known locations of their friends. Influential work in this field was conducted at Facebook Inc. in [1]. Here, the authors infer a user’s home address with a maximum likelihood estimate that best fits an empirically observed power law. Surprisingly, the results of [1] indicate that social network based methods are *more* accurate than IP-based geolocation. Since the IP addresses of Twitter users are never public, and our research involves

<sup>1</sup><http://blog.gnip.com/twitter-geo-data-enrichment>

<sup>2</sup><http://tweepstmap.com>

<sup>3</sup>[http://support.gnip.com/enrichments/profile\\_geo.html](http://support.gnip.com/enrichments/profile_geo.html)

<sup>4</sup><http://geonames.org>

<sup>5</sup>the “ego network” of a user consists of the user together with the users they are directly connected to

public data alone, we can not report any results about how the present work compares against IP-based Twitter geolocation.

Geotagging work on Flickr and Twitter by Sadilek *et al.* [25] uses social ties to infer location and then studies the converse problem: using location to infer social ties. They conclude that location alone is insufficient for this task. Very recent work by [31] uses network structure as well as language processing of user profiles to identify “landmark” users in the United States for whom location inference is optimal while [24] showcases a similar result on Korean Twitter users.

Most closely related to our work is that of [11], where the author developed a node-wise algorithm, “Spatial Label Propagation”, which iteratively estimates Twitter user locations by propagating the locations of GPS-known users across a Twitter social network. While not discussed in [11], it turns out that Spatial Label Propagation is in fact a parallel coordinate descent method applied to total variation minimization. Later in this work we show how our technique can reduce to Spatial Label Propagation by removing constraints on (3). While Spatial Label Propagation was demonstrated at scale, our study reaches higher coverage and accuracy: 101.8M users at a median error of 6.33km vs. 45.8M users at a median error of 10km.

Research indicating that Twitter contact is independent of proximity can be found in [12], where the author examines GPS-known retweet pairings and finds an average distance of 749 miles between users. Averages, however, are sensitive to outliers which are often present in social data. In this work, we will make use of robust statistics to estimate center and spread for sets of locations.

### III. METHOD

#### A. Data and Network Construction

An appropriate social network is a fundamental part of our algorithm. Twitter users often “@mention” each other by appending an “@” to the mentioned user’s name. We build a social network,  $G = (V, E)$ , with users as vertices and @mentions between users as edges.

Reciprocated @mentions indicate social ties. We define edge weights,  $w_{ij}$ , of  $G$ , the “bidirectional @mention network”, using the minimum number of reciprocated @mentions between users  $i$  and  $j$ . The key advantage to constructing a social graph from @mentions (as opposed from “followers” or “favourites”) is that it enables us to build a large social graph from a large collection of tweets without being burdened by Twitter API rate limiting.

We use a 10% sample of public tweets collected between April 2012 and April 2014. This amounts to 76.9TB of json data (uncompressed) and 25,312,399,718 @mentions. From the complete set of @mentions we built a weighted and directed network of 8,593,341,111 edges by condensing multiple mentions into weighted edges. Filtering down

to only reciprocated @mentions leaves us with a network of **1,034,362,407** edges<sup>6</sup> and **110,893,747** users. This bidirected network is the focus of our experiments.

#### B. Ground Truth User Locations

We define a function,  $f$ , which assigns to each user an estimate of their physical location. Users may opt to make their location publicly available though cellphone-enabled GPS or self-reported profile information. For this small set of users, computation of  $f$  is relatively straightforward.

To assign a unique location to a user from the set of their GPS-tagged tweets,  $\mathcal{G}$ , we compute the  $l_1$ -multivariate median [29] of the locations they have tweeted from:

$$\operatorname{argmin}_x \sum_{y \in \mathcal{G}} d(x, y) \quad (4)$$

Our data contains 13,899,315 users who have tweeted with GPS at least three times. Several of these users are highly mobile and can not reasonably be assigned a single, static, location. We filter out users whose median absolute deviation (cf. (3)) of GPS-annotated tweet locations is over 30km. This leaves with a set of 12,435,622 users whose location might be known via GPS. However, tweet timestamps reveal that 86,243 of these users have at some point exceeded the flight airspeed record of 3529.6 km/h. Manual examination of these accounts finds several bots retweeting worldwide GPS-annotated tweets as well as human users who suffer GPS malfunctions (e.g. a tweet near (0.0,0.0) shortly after a valid location). We remove from our training set any user who has travelled in excess of 1,000 km/h. The total number of GPS-known users is 12,297,785.

Following this, we extract self-reported home locations by searching through a list of 51,483 unambiguous location names for exact matches in user profiles. Users who list several locations in their profile are not geotagged by this step, though it is possible to account for such users using a method found in [31]. When self-reporting users also reveal their location though GPS, we opt to use their GPS-known location. We remove self-reports which are over 90 days old. This provides us with home locations for an additional 15,360,494 users.

The list of 51,483 location names has been optimized for accuracy on Twitter. Starting with an initial list of 67,711 location names obtained from the GeoNames project, we examined 12,471,920 GPS-tagged tweets<sup>7</sup> from users who self-reported profile locations and removed location names from the list when the median discrepancy between GPS and the reported location was greater than 100km. We have plotted the resulting discrepancies between GPS and self-reports in fig. 1.

The total number of users with ground truth locations via GPS or self-reports is **24,545,425**. Denote this set of users

<sup>6</sup>i.e. a 12% chance of @mention reciprocation

<sup>7</sup>collected between 2013-01-01 and 2013-04-07

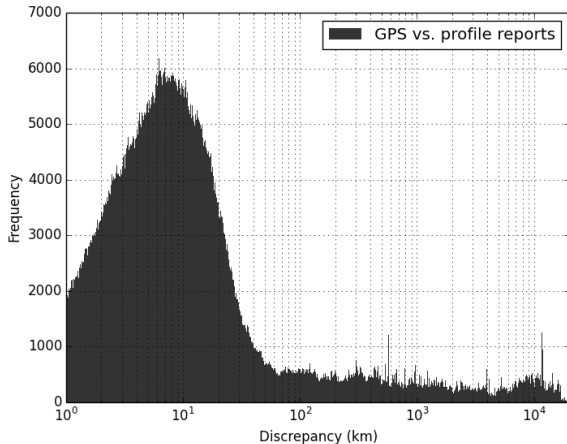


Figure 1: Histogram of discrepancies between GPS and self-reported profile locations for 1,883,331 users. The median and mean discrepancies are 7.10 and 318.50 km, respectively.

by  $L$ , and the remaining users in the network by  $U$ . The vertex set of our social network is thus partitioned as

$$V = L + U \quad (5)$$

and our goal is to assign a value of  $f$  to nodes in  $U$ .

### C. Global Optimization Algorithm

Our algorithm assigns a location to a user based on the locations of their friends. To check that online social ties are well-aligned with geographic distance, we restrict our attention to GPS-known users and study contact patterns between them in fig. 2.

Users with GPS-known locations make up only a tiny portion of our @mention networks. Despite the relatively small amount of data, we can still see in fig. 2 that online social ties typically form between users who live near each other and that a majority of GPS-known users have at least one GPS-known friend within 10km.

The optimization (1) models proximity of connected users. Unfortunately, the total variation functional is non-differentiable and finding a global minimum is thus a formidable challenge. We will employ “parallel coordinate descent” [22] to solve (1). Most variants of coordinate descent cycle through the domain sequentially, updating each variable and communicating back the result before the next variable can update. The scale of our data necessitates a parallel approach, prohibiting us from making all the communication steps required by a traditional coordinate descent method.

At each iteration, our algorithm simultaneously updates each user’s location with the  $l1$ -multivariate median of their friend’s locations. Only after all updates are complete do we communicate our results over the network.

At iteration  $k$ , denote the user estimates by  $\mathbf{f}^k$  and the variation on the  $i$ th node by

$$|\nabla_i(\mathbf{f}^k, f)| = \sum_j w_{ij} d(f, f_j^k) \quad (6)$$

Parallel coordinate descent can now be stated concisely in alg. 1.

### D. Individual Error Estimation

The vast majority of Twitter users @mention with geographically close users. However, there do exist several users who have amassed friends dispersed around the globe. For these users, our approach should not be used to infer location.

We use a robust estimate of the dispersion of each user’s friend locations to infer accuracy of our geocoding algorithm. Our estimate for the error on user  $i$  is the median absolute deviation of the inferred locations of user  $i$ ’s friends, computed via (3). With a dispersion restriction as an additional parameter,  $\gamma$ , our optimization becomes

$$\min_{\mathbf{f}} |\nabla \mathbf{f}| \text{ subject to } f_i = l_i \text{ for } i \in L \text{ and } \max_i \tilde{\nabla} f_i < \gamma \quad (7)$$

---

**Algorithm 1:** Parallel coordinate descent for dispersion-constrained TV minimization.

---

**Initialize:**  $f_i = l_i$  for  $i \in L$  and parameter  $\gamma$   
**for**  $k = 1 \dots N$  **do**  
  **parfor**  $i$  :  
    **if**  $i \in L$  **then**  
       $f_i^{k+1} = l_i$   
    **else**  
      **if**  $\tilde{\nabla} f_i \leq \gamma$  **then**  
         $f_i^{k+1} = \underset{f}{\operatorname{argmin}} |\nabla_i(\mathbf{f}^k, f)|$   
      **else**  
        no update on  $f_i$   
      **end**  
    **end parfor**  
   $\mathbf{f}^k = \mathbf{f}^{k+1}$   
**end**

---

Restricting the maximum allowed value of  $\tilde{\nabla} f_i$  during each update ensures that only reliable locations are propagated through subsequent iterations. In alg. 1 we refuse to update locations for users whose friends are dispersed beyond a given threshold.

The argument that minimizes (6) is a weighted  $l1$ -multivariate median of the locations of the neighbors of node  $i$ . By placing this computation inside the **parfor** of alg. 1 and removing any restriction on  $\gamma$ , we are able to reproduce the “Spatial Label Propagation” algorithm of [11]




Figure 2: Study of contact patterns between 953,557 users who reveal their location via GPS and are present in each of the bidirectional @mention network (red), bidirectional @mention network after filtering edges for triadic closures (green), and the unfiltered unidirectional @mention network (black).

Listing 1: Distributed implementation of alg. 1 using the Spark framework [32]

```

1 val edgeList = loadEdgeList()
2 var userLocations = loadInitialLocations()
3 for (k <- 1 to N) {
4   val adjListWithLocations = ↵
      edgeList.join(userLocations) .keyBy(x => ↵
      x._2).groupByKey()
5   val updatedLocs = adjListWithLocations.map(x ↵
      => (x._1, llMedian(x._2), dispersion(x._2)))
6   userLocations = updatedLocs.filter(x => x._3 < ↵
      GAMMA)
7 }
8 return userLocations

```

as a coordinate descent method designed to minimize total variation.

While several researchers have worked with parallel coordinate descent, existing convergence results are difficult to apply. This is in part due to the fact that most convergence studies assume Euclidean space while our algorithms involve metrics on the surface of a sphere. We also face the problem of ensuring that the ground truth users fully retain their initial location for all iterations, which prohibits us from working with an  $l_2$ -penalized unconstrained problem as is often studied in theoretical papers.

### E. Implementation Remarks

The Apache Spark cluster computing framework [32] was used to implement alg. 1. Spark allows one to distribute data in cluster memory by making use of resilient distributed datasets (referred to as RDDs) and operate on these datasets with arbitrary Scala code.

Our technique is sketched in Listing 1. Our network and user locations are stored in the RDDs `edgeList`, and `userLocations`. Computations on these RDDs make use of all available cluster cpu resources. The **parfor** of alg. 1 can be implemented with a straightforward `map` and

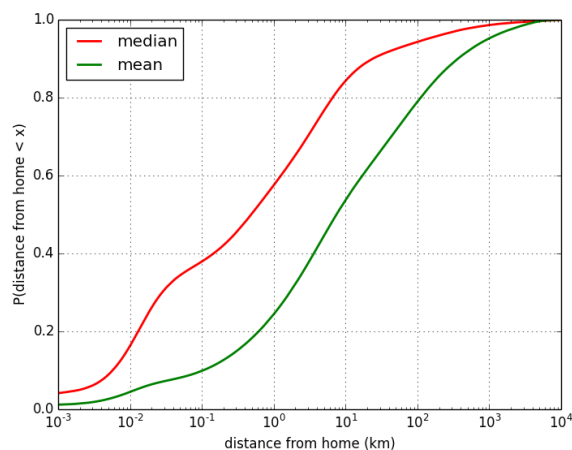


Figure 3: Empirical CDF of activity radii for GPS-known users. For each of the 13,899,315 users with 3 or more GPS-annotated tweets, we compute their home location using the median of their tweet locations and define activity radii using the mean and median distances from home.

`filter`. Communicating the updated locations across the network is accomplished with a `join` on the edge list, followed by a `groupByKey`, which sets up an adjacency list for the next `map`.

An implementation of alg. 1 taking advantage of more advanced distributed graph processing techniques, such as the Pregel model [15] or GraphX [30], is a direction for future work.

### F. Mobility Considerations

The technique outlined above is only useful for static location inference. Fast-moving users with large activity radii will be tagged incorrectly by our method. Here, we study the appropriateness of static user geotagging by restricting our

Iteration	Test users	Test users added	Median error (km)	Median error on new test users
1	771,321	771,321	5.34	5.34
2	926,019	154,698	6.02	12.31
3	956,705	30,686	6.24	45.50
4	966,515	9,810	6.32	150.60
5	971,731	5,216	6.38	232.92

Table I: Geolocated users and accuracy for each iteration. The first iterations produce the most accurate geotags and the highest coverage.

attention to GPS-annotated tweets and reporting statistics on activity radii of these users.

For the 13,899,315 users with three or more GPS-annotated tweets (cf. section III-B), we define a home location using the median of their tweet locations and plot empirical cumulative distribution functions describing the mean and median distances from home in fig. 3. The data indicates that large activity radii exist, but are atypical of Twitter users.

The presence of Twitter users with large activity radii is expected. Several recent quantitative studies of mobility patterns in public social media have confirmed that distance from home is decisively fails to follow a normal distribution. Power law, lognormal [17], gravity law [9], and radiation laws have been studied. What this means for our research is that outliers are unavoidable in geographic social data and the use of robust statistics is a must.

Large activity radii correspond to faster moving users. Timestamps on the GPS-annotated tweets reveal the median speed of a Twitter user is .02 km/h (average speed 233.63 km/h). For users with an activity radius over 20km, the median jumps an order of magnitude to 0.20 km/h (average speed 1024.04 km/h). Fast-moving accounts may be inhuman and not geotaggable. For example, the maximum speed attained by any Twitter user in our data was 67,587,505.24 km/h, slightly more than 30x the escape velocity from the surface of the Sun.

#### IV. RESULTS

We run our optimization on the bidirectional @mention network described in section III-A. Results are reported after 5 iterations of alg. 1, though high coverage can be obtained sooner (cf. table I). The parameter  $\gamma$  was set to 100km after experimenting with different values (cf. fig. 6 and fig. 8).

##### A. Coverage

To assess coverage, we examined 37,400,698,296 tweets collected between April 2012 and April 2014. These tweets were generated by 359,583,211 users.

Active Twitter users are likely present in the @mention network and are therefore likely geocoded by our algorithm. In fig. 4, we grouped users by activity level and, for each group, plotted the probability that a user is geolocated

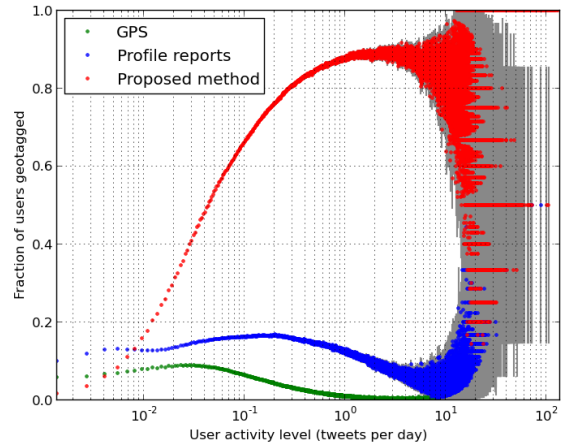


Figure 4: Fraction of geolocated users as a function of activity level. The probability that a user reveals their location via GPS (green) or self-reports (blue) changes little. The probability that a user is geotagged by our method (red) is dramatically higher for users who tweet more often. Error bars (grey) display the standard error of the mean at each activity level and are largest for high activity levels (where the number of sample points is small).

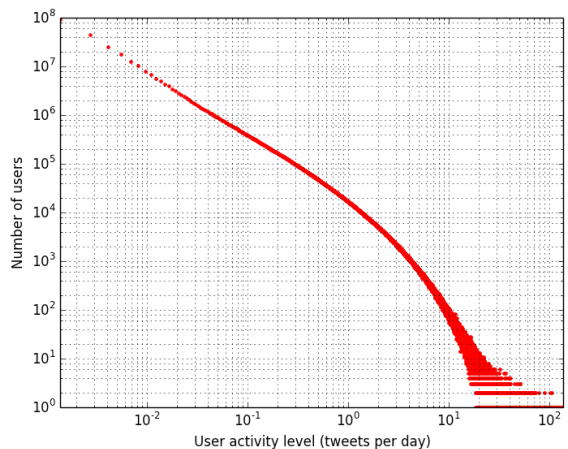


Figure 5: Log-log plot describing the total number of users as a function of activity level. A large fraction of users quickly lose interest in Twitter after generating a small number of tweets. As observed in fig. 4, these users are difficult to geocode with the proposed method.

	Tweet volume	Account volume
GPS	584,442,852 (1.6%)	12,297,785 (3.4%)
unambiguous profile	3,854,169,186 (10.3%)	45,284,996 (12.6%)
any nonempty profile	23,236,139,825 (62.1%)	164,020,169 (45.6%)
proposed method	30,617,806,498 (81.9%)	101,846,236 (28.3%)

Table II: Summary of coverage results. The sample data consisted of 37,400,698,296 tweets generated by 359,583,211 users. We are able to geotag 81.9% of tweets, more than is possible even if all nonempty self-reported profile locations were unambiguous.

by three different methods. The probability that a user is geocoded by our method increases dramatically as a function of activity level. The probability that a user is geocoded via GPS or self-reports, however, appears to be unrelated to activity level. We summarize coverage results in table II.

The statistics in fig. 4 become noisy for extremely active users. This is due to the fact that only a small number of users tweet this often (cf fig. 5). Given that we collect only 10% of Twitter, activity levels over  $10^1$  in our chart contain users who have been tweeting over 100 times per day over a two-year period which may indicate that they are not human.

### B. Accuracy

Accuracy is assessed using leave-many-out cross-validation with a 10% hold out set. From the 12,297,785 users who reveal GPS locations, we randomly selected 1,229,523 test users for exclusion from  $L$ . After 5 iterations of alg. 1 with  $\gamma = 100\text{km}$ , we were able to infer location for 971,731 test users with a median error of 6.38km and a mean error of 289.00km.

Reverse-geocoding to cities with a population over 5,000 shows that our method was accurate to city-resolution for 770,498 (89.7%) of test users. We remark here that evaluations based on semantic distance (which are common in language-based geotagging) can be difficult to compare with those based on physical distance as reverse-geocoding can introduce errors of its own. For example, depending on the convention used, a minimum population size of 5,000 would discriminate between the 48 different barrios of Buenos Aires which are physically close yet semantically distinct.

Inferring location for users with geographically dispersed ego networks leads to large errors. In fig. 6 we use data from a run with no restriction on  $\max \nabla f_i$  (a.k.a Spatial Label Propagation) and observe bimodal error distributions when restrictions on  $\max \nabla f_i$  are loose. When we account for ego network dispersion with  $\gamma$  in (7) we are able to eliminate outlying errors. The effect of modifying  $\gamma$  is further illustrated with box plots in In fig. 7.

Coverage is not substantially decreased by restricting ego network dispersion. Our experiments with tuning  $\gamma$

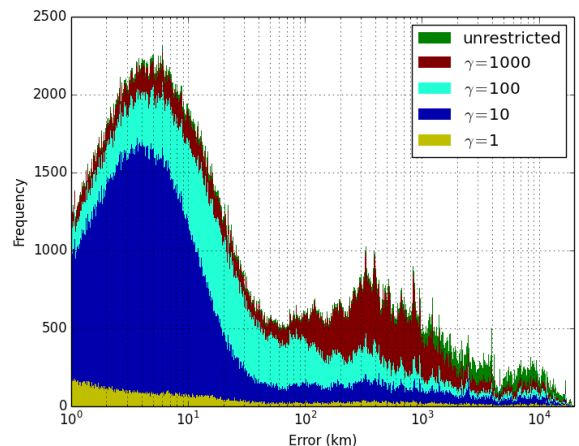


Figure 6: Histogram of errors with different restrictions on the maximum allowable geographic dispersion of each user's ego network in km (i.e.  $\gamma$  in (7)). We are able to remove outlying errors by controlling  $\gamma$ .




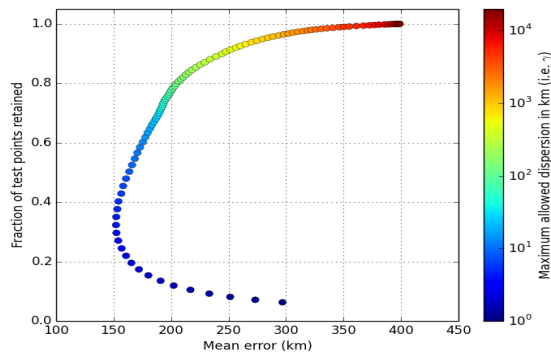
Figure 7: Box plots describing the errors for different values of  $\gamma$ . The ends of the whiskers are drawn at  $\pm 1.5$  times the interquartile range.

are visible in fig. 8 where we observe that the fraction of geolocated users remains high even when  $\gamma$  is made small.

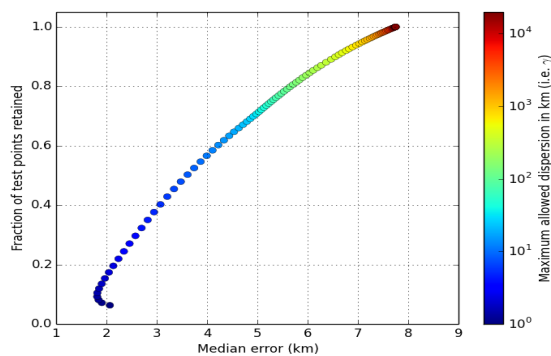
## V. CONCLUSION

We have presented a total variation-based algorithm for inferring the home locations of millions of Twitter users. By framing the social network geotagging problem as a global convex optimization, we have connected much recent work in computational social science with an immense body of existing knowledge.

Additionally, we have developed a novel technique to estimate per-user accuracy of our geotagging algorithm and used it to ensure that errors remain small.



(a) Mean error control via  $\gamma$



(b) Median error control via  $\gamma$

Figure 8: Coverage penalties and accuracy gains as a function of  $\gamma$ . In fig. 8a we observe that a restriction of  $\gamma = 100\text{km}$  reduces mean error to  $200\text{km}$  while retaining 80% of users in our geolocation database. In fig. 8b we see that median error is less responsive to  $\gamma$ , indicating that restrictions on  $\gamma$  primarily control outlying errors.

Our results have been demonstrated at scale. To the best of our knowledge, this is the largest and most accurate dataset of Twitter user locations known. The fact that it is constructable from publicly visible data opens the door for many future research directions.

#### ACKNOWLEDGMENT

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI / NBC) contract number D12PC00285. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBE, or the U.S. Government.

#### REFERENCES

- [1] Lars Backstrom, Eric Sun, and Cameron Marlow. Find me if you can: improving geographical prediction with social and spatial proximity. In *Proceedings of the 19th international conference on World wide web*, pages 61–70. ACM, 2010.
- [2] Xavier Bresson, Thomas Laurent, David Uminsky, and James von Brecht. Convergence and energy landscape for cheeger cut clustering. In *Advances in Neural Information Processing Systems 25*, pages 1394–1402, 2012.
- [3] Xavier Bresson, Thomas Laurent, David Uminsky, and James H von Brecht. Multiclass total variation clustering. In *Advances in Neural Information Processing Systems*, 2013.
- [4] Yan Chen, Jichang Zhao, Xia Hu, Xiaoming Zhang, Zhoujun Li, and Tat-Seng Chua. From interest to function: Location estimation in social media. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [5] 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.
- [6] Ryan Compton, Craig Lee, Tsai-Ching Lu, Lalindra De Silva, and Michael Macy. Detecting future social unrest in unprocessed twitter data. In *IEEE Intelligence and Security Informatics (ISI)*, 2013.
- [7] 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.
- [8] Jacob Goldenberg and Moshe Levy. Distance is not dead: Social interaction and geographical distance in the internet era. *arXiv preprint arXiv:0906.3202*, 2009.
- [9] Bartosz Hawelka, Izabela Sitko, Euro Beinat, Stanislav Sobolevsky, Pavlos Kazakopoulos, and Carlo Ratti. Geo-located twitter as the proxy for global mobility patterns. *arXiv preprint arXiv:1311.0680*, 2013.
- [10] Huiyi Hu, Thomas Laurent, Mason A Porter, and Andrea L Bertozzi. A method based on total variation for network modularity optimization using the mbo scheme. *SIAM Journal on Applied Mathematics*, 73(6):2224–2246, 2013.
- [11] David Jurgens. That’s what friends are for: Inferring location in online communities based on social relationships. *ICWSM*, 2013.
- [12] Kalev Leetaru, Shaowen Wang, Guofeng Cao, Anand Padmanabhan, and Eric Shook. Mapping the global twitter heartbeat: The geography of twitter. *First Monday*, 18(5), 2013.
- [13] Jalal Mahmud, Jeffrey Nichols, and Clemens Drews. Where is this tweet from? inferring home locations of twitter users. *Proc AAAI ICWSM*, 12, 2012.



- [14] Jalal Mahmud, Jeffrey Nichols, and Clemens Drews. Home location identification of twitter users. *ACM Transactions on Intelligent Systems and Technology*, 2014.
- [15] Grzegorz Malewicz, Matthew H Austern, Aart JC Bik, James C Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 135–146. ACM, 2010.
- [16] Benjamin Mandel, Aron Culotta, John Boulahanis, Danielle Stark, Bonnie Lewis, and Jeremy Rodrigue. A demographic analysis of online sentiment during hurricane irene. In *Proceedings of the Second Workshop on Language in Social Media*, pages 27–36. Association for Computational Linguistics, 2012.
- [17] Delia Mocanu, Andrea Baronchelli, Nicola Perra, Bruno Gonçalves, Qian Zhang, and Alessandro Vespignani. The twitter of babel: Mapping world languages through microblogging platforms. *PLoS one*, 8(4):e61981, 2013.
- [18] Delia Mocanu, Andrea Baronchelli, Nicola Perra, Bruno Gonçalves, Qian Zhang, and Alessandro Vespignani. The twitter of babel: Mapping world languages through microblogging platforms. *PLoS ONE*, 8(4):e61981, 04 2013.
- [19] Diana Mok, Barry Wellman, and Juan Carrasco. Does distance matter in the age of the internet? *Urban Studies*, 47(13):2747–2783, 2010.
- [20] Stanley Osher and Nikos Paragios. *Geometric level set methods in imaging, vision, and graphics*. Springer, 2003.
- [21] M Paul and Mark Dredze. You are what you tweet: Analyzing twitter for public health. In *Fifth International AAAI Conference on Weblogs and Social Media (ICWSM 2011)*, 2011.
- [22] Peter Richtárik and Martin Takáč. Parallel coordinate descent methods for big data optimization. *arXiv preprint arXiv:1212.0873*, 2012.
- [23] L Rudin, S Osher, and E Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, November 1992.
- [24] KyoungMin Ryoo and Sue Moon. Inferring twitter user locations with 10 km accuracy. In *Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion*, WWW Companion '14, pages 643–648, Republic and Canton of Geneva, Switzerland, 2014. International World Wide Web Conferences Steering Committee.
- [25] Adam Sadilek, Henry Kautz, and Jeffrey P Bigham. Finding your friends and following them to where you are. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 723–732. ACM, 2012.
- [26] Laura M. Smith, Matthew S. Keegan, Todd Wittman, George O. Mohler, and Andrea L. Bertozzi. Improving density estimation by incorporating spatial information. *EURASIP J. Adv. Signal Process*, 2010:7:1–7:12, February 2010.
- [27] Yuri Takhteyev, Anatoliy Gruzd, and Barry Wellman. Geography of twitter networks. *Social Networks*, 34(1):73–81, 2012.
- [28] Andranik Tumasjan, Timm O Sprenger, Philipp G Sandner, and Isabell M Welp. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the fourth international AAAI conference on weblogs and social media*, pages 178–185, 2010.
- [29] Yehuda Vardi and Cun-Hui Zhang. The multivariate 11-median and associated data depth. *Proceedings of the National Academy of Sciences*, 97(4):1423–1426, 2000.
- [30] Reynold S Xin, Joseph E Gonzalez, Michael J Franklin, Ion Stoica, and EECS AMPLab. Graphx: A resilient distributed graph system on spark. *GRADES (SIGMOD workshop)*, 2013.
- [31] Yuto Yamaguchi, Toshiyuki Amagasa, and Hiroyuki Kitagawa. Landmark-based user location inference in social media. In *Proceedings of the First ACM Conference on Online Social Networks*, COSN '13, pages 223–234, New York, NY, USA, 2013. ACM.
- [32] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael Franklin, Scott Shenker, and Ion Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 2–2. USENIX Association, 2012.