



Time-Critical Geolocation for Social Good

Reem Suwaileh (✉)

Qatar University, Doha, Qatar

rs081123@qu.edu.qa

Abstract. Twitter has become an instrumental source of news in emergencies where efficient access, dissemination of information, and immediate reactions are critical. Nevertheless, due to several challenges, the current fully-automated processing methods are not yet mature enough for deployment in real scenarios. In this dissertation, I focus on tackling the lack of context problem by studying automatic geo-location techniques. I specifically aim to study the *Location Mention Prediction* problem in which the system has to extract location mentions in tweets and pin them on the map. To address this problem, I aim to exploit different techniques such as training neural models, enriching the tweet representation, and studying methods to mitigate the lack of labeled data. I anticipate many downstream applications for the Location Mention Prediction problem such as incident detection, real-time action management during emergencies, and fake news and rumor detection among others.

Keywords: Geolocation · Social good · Twitter

1 Introduction

Twitter, as a cost-effective and time-saving communication channel, plays an important role in emergencies. It becomes a substantial information source for response and recovery [28] and for grassroots management in relief activities [10, 27]. Invaluable efforts have been made to utilize social media for preparedness, relief, and recovery during natural disasters and after they are over [11]. However, the current automatic solutions require human intervention in many stages due to many challenges. Among these challenges is the lack of geographical context that is important for response and relief. For example, during a disaster, responders need to locate the incidents (e.g., road closures, infrastructure damage, etc.) as they happen, tweets, and users discussing the disaster. However, people often tend to hide their geographical information due to privacy and safety concerns [16, 22]. Anderson et al. [12] analysed tweet disaster datasets that span over a period of 6 years and showed only around 2% or less of the tweets are geo-referenced. Thus, developing automatic geolocation tools would enable real-time location-aware monitoring of the disaster which makes the decision-making process more reliable, effective, and efficient.

In my work, I am interested in tackling the *Location Mention Prediction* (LMP) problem during time-critical situations. The problem involves two tasks

that can be tackled separately or jointly: (1) *Location Recognition*: extracting the location mentions in tweets, and (2) *Location Disambiguation*: locating potential location mentions on the map. The *Disambiguation* task includes two identification sub-tasks: (2.1.) identifying the *intended location* from a set of location mentions sharing the same toponym. (2.2.) identifying the *locational focus* of a tweet containing different location mentions. Learning to predict the location mentions is a non-trivial task. The location taggers have to address many challenges including microblogging-specific challenges (e.g., tweet sparsity, noisiness, stream rapid-changing, hashtag riding, etc.) and the task-specific challenges (e.g., time-criticality of the solution, scarcity of labeled data, etc.). While tackling these challenges, I aim to address several research questions including: **RQ1**. Are deep learning approaches more effective compared to the state-of-the-art and traditional machine learning-based LMP approaches?, **RQ2**. Would context expansion (using user’s tweets, on-topic tweets, etc) improve LMP?, **RQ3**. How can we reduce the effect of the scarcity of labeled data on the performance of the LMP system?, and **RQ4**. How can the LMP systems control the trade-off between effectiveness and efficiency during crisis scenarios?.

The remainder of this paper is organized as follows. The related literature is reviewed in Sect. 2. My proposed methodology is discussed in Sect. 3 followed by the evaluation setup in Sect. 4.

2 Related Work

In this section, I discuss the related work to LMP problem over tweets.

TwitterStand [23] is a tweet geo-tagging system for extracting breaking news and pinning them on the map. Lingad et al. [18] compared a few NER tools on disaster-related Twitter data and found their performance noticeably degraded over the Twitter stream. Li et al. [14], on the other hand, constructed their own noisy gazetteer using a crowdsourcing-like method to match extracted location mentions from tweets by the POI tagger. Malmasi et al. [19] extracted noun phrases (NPs) in tweets using a recursive rule-based tree parser and link potential locations with Geonames entries using fuzzy matching. Ghahremanlou et al. [7] explored combined techniques to identify the location mentions by both matching and StanfordNER. The major weakness in gazetteer-based methods is the mismatch between the noisy Twitter stream and non-noisy gazetteer entries [17]. To address this issue, Li et al. [15], constructed their own noisy gazetteer using collected cross-posts on Twitter from Foursquare check-ins. Alternatively, Sultanik and Fink [25], used Information Retrieval (IR) based approach to identify the location mentions in tweets. Unlike Ghahremanlou et al. [7], Yin et al. [29] retrained StanfordNER using tweet dataset to effectively identify the location mentions in tweets. More interestingly, to achieve high coverage of recognized locations, a couple of studies [6, 30] adopt an ensemble-based parser.

In 2014, the topic of the fifth Australasian Language Technology Association (ALTA) shared task was on identifying location mentions in tweets [21]. Participants explored several techniques such as feature engineering, ensemble

classifiers, rule-based classification, knowledge infusion, CRFs sequence labelers, semi-supervision. Al-Olimat et al. [1] proposed identifying the location names by traversing a tree of the tweet’s n-grams to extract valid locations that exist in their pre-build region-specific gazetteer. Moreover, Hoang and Mothe [8] combined syntactic and semantic features to train traditional ML-based models whereas Kumar and Singh [13] trained a Convolutional Neural Network (CNN) model that learns the continuous representation of tweet text and then identifies the location mentions.

The gap in existing solutions is two-fold. First, in relation to methods, a few studies investigated deep learning-based solutions, most of the proposed solutions are gazetteer-based, and most of them do not consider efficiency when developed. Second, there is not a unified evaluation framework in which a few small-scale datasets are available and different tools are compared. Additionally, the efficiency of the proposed methods is rarely evaluated.

3 Proposed Research

In this section, I describe the proposed solutions to address the research questions listed in Sect. 1.

Deep Location Prediction (RQ1): I perceive the location recognition task as a multi-label classification task. I opt to use the Neural Networks (NNs) algorithms due to their ability to learn features and model parameters simultaneously from incomplete or noisy training data [4]. I specifically plan to experiment with (1) the Bidirectional Long-Term Short-Term Memory [9, 24], (2) Encoder-Decoder with attention [2, 26], and (3) BERT with Fine-tuning [5]. For the disambiguation task, I plan to explore the effectiveness of Siamese Neural Networks [3] that was used recently in neural-based IR models, especially for short text matching [20]. I further plan to experiment with character n-grams to better capture the lexical information.

Context Expansion (RQ2): Due to the short length of tweets, systems lack the context that would enable them to detect the location mentions effectively. To enrich the context of tweet, I plan to explore four tweet expansion sources including (1) User’s tweets: I hypothesize that tweets shared by the user within a time window, say 10 min, are most probably discussing the same topic, (2) On-topic tweets: I assume that tweets sharing the same trending related hashtag to the disaster to be topically relevant, (3) Linked webpages: I anticipate the URLs to be useful sources for enriching the tweet context, and (4) Knowledge-bases (KB): I hypothesize that entity recognition and analysis using external auxiliary data, e.g., knowledge-bases, can aid understanding the spatial focus of a tweet which in turn enables LMP. I plan to use general-purpose KBs and study the effectiveness of knowledge-base population and acceleration techniques to maintain an online up-to-date KBs during the disaster.

Handling Data Scarcity (RQ3): Deep learning algorithms are data-hungry which requires budget and expensive resources for acquiring labeled data. To address this challenge, I investigate possible ways to reduce its effect during

disasters such as (1) Exploiting existing data: Using training data from past disasters of similar or a different disaster type to train prediction systems. I plan to leverage one-step domain adaptation techniques (e.g., divergence-based, etc.), (2) Acquiring cheaper data: I plan to explore the effectiveness of expanding the small labeled datasets of a current disaster using semi-supervision, weak supervision and active learning methods, and (3) Reusing pre-trained tools: I plan to study the effectiveness of already-trained tools (based on their availability) on old disasters for effective LMP on new disasters (e.g., transfer learning).

Effectiveness and Efficiency Trade-off (RQ4): As I plan to tackle the recognition task in the disaster domain, I aim to train my models in *real-time* while the disaster is happening. Thus, I plan to investigate the trade-off between effectiveness and efficiency of LMP systems. Possible paths to study this trade-off are: (1) Tuning system decision by, for example, prioritizing tweets for LMP instead of checking every tweet chronologically, (2) Analyzing time and space complexities, and (3) exploring possible ways to modify the effectiveness measures to account for efficiency.

4 Experimental Evaluation

To evaluate the LMP proposed approaches, precision, recall, and F1 scores will be computed. When a system manages to extract part of the location mention, it is penalized by counting the false positives and false negatives multiplied by the percentage of overlap between the system’s output and the ground-truth. To conduct the initial evaluation, the publicly-available LMP English datasets, that are samples of disaster-specific streams [1, 21, 29], will be used. I anticipate the solutions to generalize to other data domains sharing the same properties with the Twitter stream. The Geonames¹ and OpenStreetMap² gazetteers will be utilized for the evaluation of the *Disambiguation task*. The approaches reviewed in Sect. 2, according to their availability and reproducibility, are the baselines against the proposed approaches for all tasks: (1) Twitter-based location mention detection and disambiguation tool (e.g., LNE_x [1]), (2) Academic NER (e.g., StanfordNER³, etc.), and (3) Commercial NER taggers (e.g., Google NL⁴, etc.).

Acknowledgments. This work was made possible by GSRA grant# GSRA5-1-0527-18082 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

References

1. Al-Olimat, H.S., Thirunarayan, K., Shalin, V., Sheth, A.: Location name extraction from targeted text streams using gazetteer-based statistical language models. In: COLING, pp. 1986–1997 (2018)

¹ <http://www.geonames.org/>.

² <http://www.openstreetmap.org/>.

³ <http://nlp.stanford.edu/software/CRF-NER.shtml>.

⁴ <https://cloud.google.com/natural-language/>.

2. Cho, K., et al.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint [arXiv:1406.1078](https://arxiv.org/abs/1406.1078) (2014)
3. Chopra, S., Hadsell, R., LeCun, Y., et al.: Learning a similarity metric discriminatively, with application to face verification. In: CVPR, vol. 1, pp. 539–546 (2005)
4. Deroncourt, F., Lee, J.Y., Szolovits, P.: NeuroNER: an easy-to-use program for named-entity recognition based on neural networks. In: EMNLP (2017)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: pre-training of deep bidirectional transformers for language understanding. arXiv (2018)
6. Gelernter, J., Balaji, S.: An algorithm for local geoparsing of microtext. *GeoInformatica* **17**(4), 635–667 (2013)
7. Ghahremanlou, L., Sherchan, W., Thom, J.A.: Geotagging twitter messages in crisis management. *Comput. J.* **58**(9), 1937–1954 (2014)
8. Hoang, T.B.N., Mothe, J.: Location extraction from tweets. *Inf. Process. Manag.* **54**(2), 129–144 (2018)
9. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
10. Hughes, A.L., Palen, L.: Twitter adoption and use in mass convergence and emergency events. *IJEM* **6**, 248–260 (2009)
11. Imran, M., Castillo, C., Diaz, F., Vieweg, S.: Processing social media messages in mass emergency: a survey. *ACM Comput. Surv.* **47**(4), 1–38 (2014)
12. Jurgens, D., Finethy, T., McCorriston, J., Xu, Y.T., Ruths, D.: Geolocation prediction in twitter using social networks: a critical analysis and review of current practice. In: ICWSM (2015)
13. Kumar, A., Singh, J.P.: Location reference identification from tweets during emergencies: a deep learning approach. *IJDRR* **33**, 365–375 (2019)
14. Li, C., Sun, A.: Fine-grained location extraction from tweets with temporal awareness. In: SIGIR, pp. 43–52 (2014)
15. Li, C., Sun, A.: Extracting fine-grained location with temporal awareness in tweets: a two-stage approach. *JASIST* **68**(7), 1652–1670 (2017)
16. Li, X., Caragea, D., Zhang, H., Imran, M.: Localizing and quantifying damage in social media images. In: IEEE/ACM ASONAM, pp. 194–201. IEEE (2018)
17. Lieberman, M.D., Samet, H., Sankaranarayanan, J.: Geotagging with local lexicons to build indexes for textually-specified spatial data. In: ICDE, pp. 201–212 (2010)
18. Lingad, J., Karimi, S., Yin, J.: Location extraction from disaster-related microblogs. In: Proceedings of the 22nd International Conference on World Wide Web, pp. 1017–1020. ACM (2013)
19. Malmasi, S., Dras, M.: Location mention detection in tweets and microblogs. In: PACLING, pp. 123–134 (2015)
20. Mitra, B., Craswell, N., et al.: An introduction to neural information retrieval. *Found. Trends® Inf. Ret.* **13**(1), 1–126 (2018)
21. Molla, D., Karimi, S.: Overview of the 2014 alta shared task: identifying expressions of locations in tweets. In: ALTA Workshop, pp. 151–156 (2014)
22. Murdock, V.: Your mileage may vary: on the limits of social media. *SIGSPATIAL Special* **3**(2), 62–66 (2011)
23. Sankaranarayanan, J., Samet, H., Teitler, B.E., Lieberman, M.D., Sperling, J.: Twitterstand: news in tweets. In: ACM SIGSPATIAL, pp. 42–51 (2009)
24. Schuster, M., Paliwal, K.K.: Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.* **45**(11), 2673–2681 (1997)
25. Sultanik, E.A., Fink, C.: Rapid geotagging and disambiguation of social media text via an indexed gazetteer. *Proc. ISCRAM* **12**, 1–10 (2012)

26. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: *Advances in neural information processing systems*, pp. 3104–3112 (2014)
27. Veil, S.R., Buehner, T., Palenchar, M.J.: A work-in-process literature review: Incorporating social media in risk and crisis communication. *J. Contingencies Crisis Manag.* **19**(2), 110–122 (2011)
28. Vieweg, S.E.: Situational awareness in mass emergency: a behavioral and linguistic analysis of microblogged communications. Ph.D. thesis (2012)
29. Yin, J., Karimi, S., Lingad, J.: Pinpointing locational focus in microblogs. In: *ADCS*, p. 66 (2014)
30. Zhang, W., Gelernter, J.: Geocoding location expressions in twitter messages: a preference learning method. *JOSIS* **9**, 37–70 (2014)