

‘Watch the Flu’: A Tweet Monitoring Tool for Epidemic Intelligence of Influenza in Australia

Brian Jin,^{*♣} Aditya Joshi,^{*♣} Ross Sparks,[♣] Stephen Wan,[♣] Cécile Paris,[♣] C Raina MacIntyre[♣]

[♣]CSIRO Data61, Sydney, Australia

[♣]The Kirby Institute, University of New South Wales, Sydney, Australia
 {firstname.lastname}@csiro.au, r.macintyre@unsw.edu.au

Abstract

‘Watch The Flu’ is a tool that monitors tweets posted in Australia for symptoms of influenza. The tool is a unique combination of two areas of artificial intelligence: natural language processing and time series monitoring, in order to assist public health surveillance. Using a real-time data pipeline, it deploys a web-based dashboard for visual analysis, and sends out emails to a set of users when an outbreak is detected. We expect that the tool will assist public health experts with their decision-making for disease outbreaks, by providing them insights from social media.

Epidemic intelligence is an area of public health surveillance that deals with early identification and assessment of potential health hazards (Paquet et al. 2006). This could be done using official information sources such as health records, or unofficial sources such as social media (Joshi et al. 2019b). In this demonstration, we present the first version of ‘Watch The Flu’, a near-real-time tool for social media-based epidemic intelligence. It currently monitors tweets posted in Australia for symptoms of influenza. Influenza is a syndrome (*i.e.*, a collection of symptoms) that results in an acute respiratory infection (<https://www.who.int/influenza/en/>). We monitor seven symptoms of influenza (namely, cold, vomiting, upset stomach, cough, headache, fever and diarrhoea) over six states of Australia (New South Wales, Victoria, Queensland, Western Australia, South Australia and Tasmania). Past work using search query statistics to predict influenza infection counts has been criticised for over-estimating the counts (Lazer et al. 2014). Therefore, our tool does not predict infection counts. Instead, it endeavours to signal outbreaks of individual symptoms as indicated by an unexpectedly high number of tweets reporting them. ‘Watch the Flu’ is expected to assist health professionals and public health experts by providing them insights from tweets in a region. The tool generates two kinds of outputs: emails that are sent when an alert is flagged and a dashboard that can be consulted on an ongoing basis.

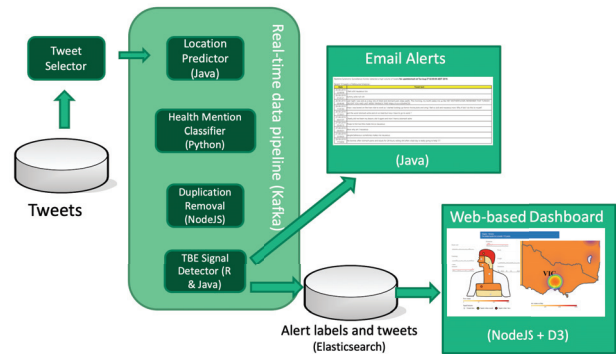


Figure 1: Architecture




Figure 2: Screenshot of the Web-based Dashboard.

Architecture

The architecture is shown in Figure 1 and described hereafter. The tool reads from a real-time database of tweets posted in Australia (‘tweet selector’), and performs a four step process to detect the alerts. The first step (‘location predictor’) assigns a tweet to one of the six states of Australia while the latter three are based on our architecture for social media-based epidemic intelligence (Joshi et al. 2019c). Programming languages used to implement each step are indicated in brackets in the Figure. ElasticSearch, Apache Kafka and NodeJS run on Docker.

Tweet selectors: Tweet selector uses two systems that collect tweets posted from Australia in real time: the Vizie (Wan, Paris, and Georgakopoulos 2015) and ESA (Cameron et al. 2012). Keywords related to the symptoms as given in (Sparks et al. 2017) are used to select relevant tweets.

Location predictor (Java): This step predicts the Australian state from which a tweet could have been posted, using a combination of heuristics based on metadata, a label propagation method, timezone text classifiers, and an information retrieval approach. (Jayasinghe et al. 2016)

Health mention classifier (Python): Mention of a symptom word in a tweet does not mean that it reports the symptom. Therefore, the health mention classifier predicts whether or not a tweet refers to a health condition. We deploy our best performing health mention classifier as tested on a held-out dataset. It uses pre-trained language models from Universal Sentence Encoder (Cer et al. 2018) as the feature representation, with SVM-Perf (Joachims 2005) optimised for F-score (Joshi et al. 2019a). The classifier is trained on benchmark datasets (Jiang, Calix, and Gupta 2016; Karisani and Agichtein 2018).

Duplication removal (NodeJS): For every user, only the first health mention tweet posted on a given day is retained. In (Joshi et al. 2019c), we describe that this step allows as a guard against reports of one person flooding the monitoring algorithm.

Time-between-events (TBE) Signal Detector (R & Java): TBE is the duration between consecutive tweets that have been detected as health mentions. A short TBE indicates that health mention tweets are being posted frequently, implying a likely alert. Using the algorithm given in (Sparks et al. 2019), we implement a TBE signal detector that uses expected and observed values of TBE calculated on the publish timestamp of health mention tweets. This algorithm runs separately for each symptom-state combination.

Output: The system generates two kinds of outputs. When an alert for a symptom-state combination is detected, an email is sent to a pre-determined set of email addresses using the JavaMail API. The email contains 10 most recent tweets for the combination. The second output is a web-based dashboard as shown in Figure 2. Symptoms are represented on parts of a human body illustration while the states are on a map of Australia obtained from the Australian Bureau of Statistics. Tweets on the map are visualised in the form of a heat map using Simpleheat (<http://mourner.github.io/simpleheat/demo/>). For each state, counts of tweets associated with the symptom are represented as line graphs next to the corresponding parts of the human body illustration. For example, the red colour in the head and the line graph next to it indicate a likely outbreak of headache reports in the state of Victoria (indicated in the map). The tool loops over different states every few seconds to display the symptom outbreaks statistics for the state.

Summary

‘Watch the Flu’ is a tweet monitoring tool that can supplement the expertise of epidemiologists and public health experts with insights from publicly available tweets. It is cur-

rently restricted to a specific syndrome (influenza) and geographical location (Australia). In future versions of ‘Watch The Flu’, we hope to include a larger set of data sources, a wider set of syndromes and be able to signal outbreaks in terms of a collection of symptoms.

References

- Cameron, M. A.; Power, R.; Robinson, B.; and Yin, J. 2012. Emergency situation awareness from twitter for crisis management. In *Proceedings of the 21st International Conference on World Wide Web*, 695–698.
- Cer, D.; Yang, Y.; Kong, S.-y.; Hua, N.; Limtiaco, N.; et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Jayasinghe, G.; Jin, B.; McHugh, J.; Robinson, B.; and Wan, S. 2016. CSIRO Data61 at the WNUT geo shared task. In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, 218–226.
- Jiang, K.; Calix, R.; and Gupta, M. 2016. Construction of a personal experience tweet corpus for health surveillance. In *Proceedings of the 15th workshop on biomedical natural language processing*, 128–135.
- Joachims, T. 2005. A support vector method for multivariate performance measures. In *Proceedings of the 22nd International conference on Machine learning*, 377–384.
- Joshi, A.; Karimi, S.; Sparks, R.; Paris, C.; and MacIntyre, C. R. 2019a. A comparison of word-based and context-based representations for classification problems in health informatics. In *Proceedings of the 18th BioNLP Workshop*.
- Joshi, A.; Karimi, S.; Sparks, R.; Paris, C.; and MacIntyre, C. R. 2019b. Survey of text-based epidemic intelligence: A computational linguistic perspective. *ACM Computing Surveys (In Press)*.
- Joshi, A.; Sparks, R.; McHugh, J.; Karimi, S.; Paris, C.; and MacIntyre, C. R. 2019c. Harnessing tweets for early detection of an acute disease event. *Epidemiology (In Press)*.
- Karisani, P., and Agichtein, E. 2018. Did you really just have a heart attack?: towards robust detection of personal health mentions in social media. In *Proceedings of the 2018 World Wide Web Conference*, 137–146.
- Lazer, D.; Kennedy, R.; King, G.; and Vespignani, A. 2014. The parable of google flu: traps in big data analysis. *Science* 343(6176):1203–1205.
- Paquet, C.; Coulombier, D.; Kaiser, R.; and Ciotti, M. 2006. Epidemic intelligence: a new framework for strengthening disease surveillance in europe. *Euro Surveill* 11:212–214.
- Sparks, R. S.; Robinson, B.; Power, R.; Cameron, M.; and Woolford, S. 2017. An investigation into social media syndromic monitoring. *Communications in Statistics - Simulation and Computation* 46(8):5901–5923.
- Sparks, R.; Jin, B.; Karimi, S.; Paris, C.; and MacIntyre, C. R. 2019. Real-time monitoring of events applied to syndromic surveillance. *Quality Engineering* 31(1):73–90.
- Wan, S.; Paris, C.; and Georgakopoulos, D. 2015. Improving government services using social media feedback. In *Social Media for Government Services*. Springer. 221–246.