

Forecasting Event-Driven Connotation Dynamics and Classifying Event Types on Twitter with LSTMs

Hannah Rashkin, Yejin Choi, and Svitlana Volkova

The types of events reported in real-time on social media (e.g. Twitter) and the connotation (e.g., perspectives) with which they are discussed are interconnected, dynamic phenomena. When a major world event occurs, topics of discussion change as do the attitudes towards the people and organizations involved. Moreover, these changes can vary depending on a Twitter user’s background, such as where they live or what language they speak. To explore this phenomenon, we have created and analyzed a dataset of 6.3 million tweets produced by news accounts in 50 countries and 11 different languages taken from a two-week period around the Brussels attacks on March 22, 2016. We parsed these tweets and extracted part-of-speech tags using SyntaxNet¹.

This work focuses on three main contributions: (1) forecasting entity and event-driven connotation dynamics, (2) event type classification using neural network models, and (3) analyzing variations and biases in Twitter news accounts perspectives towards salient entities (e.g., people, organizations) across three dimensions – time, language, and location.

First, we trained Long Short-term Memory (LSTMs) models to predict the distribution of perspectives towards an entity on a given day based on the previous days. Other work has investigated the dynamics of overt sentiment and events using time-series analysis [7]. We add to this body of work by using deep learning to analyze connoted perspective, a subtler aspect of sentiment. We run experiments for languages other than English while tracking perspective per country, making this one of the first works to detect changes in sentiment over time with separation of location and language. Analysis of the results show that an LSTM was able to perform well at this task (KL divergence of 1.8 and pearson correlation of 0.8 in English) as well as the harder task of predicting perspective much later in time.

Next, we built upon previous work in event detection by creating our own Twitter-specific heuristic for event clustering [1, 6, 3]. Our event clusters are used to automatically predict the event types indicated by a given tweet (e.g., attack, natural disaster), inspired by papers that have previously codified event types [4, 2]. We compared the performance on this task of popular models such as Logistic Regression vs. RNN models with distributed word representations [5], with our results showing the best performance in LSTM models. These results were extended to the task of modeling event types sequentially to predict what types of events will be reported tomorrow.

Finally, we performed an exploratory analysis of Twitter news accounts perspectives towards salient entities (e.g. people, organizations). We computed these perspectives using the connotation frame lexicon², which captures nuanced connotative relationships (e.g., writer’s perspective towards agent or theme). One contribution of our work is that we have extended this resource to be

¹<https://github.com/tensorflow/models/tree/master/syntaxnet>

²<http://homes.cs.washington.edu/hrashkin/connframe.html>

usable with 10 non-English languages by propagating to other languages via parallel corpora. Our analysis is one of the first to compare perspective across three dimensions – time, language, and location. We demonstrate empirically that the general public’s perspective changes over time and also varies across linguistic and geographic borders.

References

- [1] Farzindar Atefeh and Wael Khreich. A survey of techniques for event detection in Twitter. *Computational Intelligence*, 2015.
- [2] Ann Bies, Zhiyi Song, Jeremy Getman, Joe Ellis, Justin Mott, Stephanie Strassel, Martha Palmer, Teruko Mitamura, Marjorie Freedman, Heng Ji, et al. A comparison of event representations in deft.
- [3] Andrew J McMinn, Yashar Moshfeghi, and Joemon M Jose. Building a large-scale corpus for evaluating event detection on twitter. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 409–418. ACM, 2013.
- [4] Donald Metzler, Congxing Cai, and Eduard Hovy. Structured event retrieval over microblog archives. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 646–655. Association for Computational Linguistics, 2012.
- [5] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *Proceedings of Workshop at ICLR*, 2013.
- [6] Alan Ritter, Oren Etzioni, Sam Clark, et al. Open domain event extraction from twitter. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1104–1112. ACM, 2012.
- [7] Mikalai Tsytsarau, Themis Palpanas, and Malu Castellanos. Dynamics of news events and social media reaction. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 901–910. ACM, 2014.