Leveraging the Power of Social Media: Talk Abstract

Leon Derczynski
Natural Language Processing Group
Department of Computer Science
University of Sheffield
leon@dcs.shef.ac.uk

Abstract

This abstract contains a summary and further reading to accompany the talk "Leveraging the Power of Social Media", given at the second University of Sheffield Engineering Symposium in the AI and VR track, on 24^{th} June 2014.

1 Introduction

From a business and government point of view there is an increasing need to interpret and act upon information from large-volume, social media streams, such as Twitter, Facebook, and forum posts. While natural language processing from newswire has been very well studied in the past two decades, understanding social media content has only recently been addressed in NLP research.

Social media poses three major computational challenges, dubbed by Gartner the 3Vs of big data: volume, velocity, and variety. NLP methods, in particular, face further difficulties arising from the short, noisy, and strongly contextualised nature of social media. To address the 3Vs of social media, novel language technologies have emerged, e.g. using locality sensitive hashing to detect new stories in media streams (volume), predicting stock market movements from tweet sentiment (velocity), and recommending blogs and news articles based on users' own comments (variety).

2 Social Media

Social media can be cast as big data, exhibiting Gartner's 3V definition: having velocity, volume and variety. There are many networks that we could study. We choose Twitter as a model organism for social media (Tufekci, 2014).

3 Understanding

Doing NLP over social media presents challenges. The "non-standard" use of language in online social networks has sometimes been perceived as due to mistakes, but in fact is often highly structured (Hu et al., 2013). For example, g-dropping is consistently mapped from speech to online language use (Eisenstein et al., Proceedings of the Annual Meeting of the Association for Computational Linguistics). Literacy scores exhibit no significant difference between users of digital slang and those who communicate using standard text (Drouin and Davis, 2009). Even mis-spellings can be deliberate and indicate social variables; for example, the use of "luccy" over "lucky" suggests that the author may be in a gang affiliated with the crips, who refuses to type "ck" which can stand for "crip killer" (Stewart, 2014). Standardising text and removing this information is lossy. We face particular challenges in identifying the language used in tweets (Lui and Baldwin, 2012; Derczynski et al., 2013a); in correctly separating words on social networks (Derczynski et al., 2013a); and in finding names (Derczynski et al., 2014), of e.g. people in tweets (Derczynski and Bontcheva, 2014c).

4 Old News

There exist strong examples of twitter improving life quality. As suggested by the webcomic XKCD,[1] tweets can outrun an earthquake after the initial delay in sending the message. This has led to a significant

[1]http://xkcd.com/723/

reduction – in the order of a few seconds – in both US and Japanese quake early warning and response systems, critical for people escaping buildings (Crooks et al., 2012; Sakaki et al., 2010).

We know that West Nile Virus affects crows before humans. Monitoring spatio-temporal spikes in animal deaths (Xu et al., 2012) allows hospitals to prepare capacity and take care of vulnerable humans in advance of outbreaks (Sugumaran and Voss, 2012), at very low recurring cost.

Based on detecting whether a person is ill based on self-reporting in tweets allows us, through the person's social network and GPS usage, to predict over an entire city whether a given twitterer will contract flu or a cold within the next eight days with about 80% accuracy (Sadilek et al., 2012).

Finally, using tools built at Sheffield (Bontcheva et al., 2013; Derczynski et al., 2013b), Australian crisis responses are integrating tweets as a sensor for fire engine dispatch to bushfire incidences (Power et al., 2013).

5 Trust

These critical uses of social media make trust important in social networks; we need to know who can be trusted, and which messages are reliable. The PHEME project (EU, 36-month) is aimed at addressing this (Derczynski and Bontcheva, 2014b; Derczynski and Bontcheva, 2014a), finding false claims and classifying them as rumours, speculation, misinformation or disinformation.

It is also important to be responsible about data use; retaining private records for too long, or retaining too much, is harmful. Cutting down data retention to just metadata is not enough (Opsahl, 2013).

6 Bias

Why is social media so hard to process with existing NLP tools? Many tools have been developed using just newswire text, which has some very specific biases (Eisenstein, 2013). Twitter includes many different kinds of user, each with their own style (Hu et al., 2013). Escaping the singular news style which biases so many of our corpora and instead examining a broad range of text styles causes problems (Baldwin et al., 2013), or rather, highlights deficiencies in our existing NLP toolkit.

7 Conclusion

We're racing ahead and improving life quality directly through social media. The research at Sheffield provides tools for dealing with this kind of text, which can help any application that relies on text from social networks. There is immense value in this social network "trivia". As a direct result, understanding social media lets us help people better.

References

T. Baldwin, P. Cook, M. Lui, A. MacKinlay, and L. Wang. 2013. How noisy social media text, how diffrnt social media sources. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 356–364. ACL.

K. Bontcheva, L. Derczynski, A. Funk, M. A. Greenwood, D. Maynard, and N. Aswani. 2013. TwitIE: An Open-Source Information Extraction Pipeline for Microblog Text. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. Association for Computational Linguistics.

Crooks, Croitoru, Stefanidis, and Radzikowski. 2012. #Earthquake: Twitter as a Distributed Sensor System. *Transactions on Geographic Information Systems*, 17(1):124.

L. Derczynski and K. Bontcheva. 2014a. Pheme: Veracity in Digital Social Networks. In *Proceedings of the UMAP Project Synergy workshop*.

L. Derczynski and K. Bontcheva. 2014b. Spatio-temporal grounding of claims made on the web, in Pheme. In *Proceedings of the 10th joint ACL-ISO workshop on Interoperable Semantic Annotation*. ACL.

L. Derczynski and K. Bontcheva. 2014c. Passive-aggressive sequence labeling with discriminative post-editing for recognising person entities in tweets. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, volume 2, pages 69–73.

L. Derczynski, D. Maynard, N. Aswani, and K. Bontcheva. 2013a. Microblog-Genre Noise and Impact on Semantic Annotation Accuracy. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*. ACM.

L. Derczynski, A. Ritter, S. Clark, and K. Bontcheva. 2013b. Twitter Part-of-Speech Tagging for All: Overcoming Sparse and Noisy Data. In *Proceedings of Recent Advances in Natural Language Processing (RANLP)*. Association for Computational Linguistics.

L. Derczynski, D. Maynard, G. Rizzo, M. van Erp, G. Gorrell, R. Troncy, and K. Bontcheva. 2014. Analysis of named entity recognition and linking for tweets. *Information Processing and Management*, To appear.

M. Drouin and C. Davis. 2009. R u txting? is the use of text speak hurting your literacy? *Journal of Literacy Research*, 41(1):46–67.

J. Eisenstein, B. O'Connor, N. Smith, and E. Xing. Proceedings of the Annual Meeting of the Association for Computational Linguistics. Discovering sociolinguistic associations with structured sparsity. In *2010*.

J. Eisenstein. 2013. What to do about bad language on the internet. In *Proceedings of NAACL-HLT*, pages 359–369.

Y. Hu, K. Talamadupula, S. Kambhampati, et al. 2013. Dude, srsly?: The surprisingly formal nature of twitters language. *Proceedings of ICWSM*.

M. Lui and T. Baldwin. 2012. langid. py: An off-the-shelf language identification tool. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL 2012), Demo Session, Jeju, Republic of Korea*.

K. Opsahl. 2013. Everything We Know About NSA Spying: "Through a PRISM, Darkly". In *Proceedings of the Chaos Communication Congress*.

Power, Robinson, and Ratcliffe. 2013. Finding Fires with Twitter. In *Proceedings of the Australasian Language Technology Association Workshop*.

A. Sadilek, H. Kautz, and V. Silenzio. 2012. Modeling spread of disease from social interactions. In *International AAAI Conference on Weblogs and Social Media (ICWSM)*, pages 322–329. AAAI.

T. Sakaki, M. Okazaki, and Y. Matsuo. 2010. Earthquake shakes Twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World Wide Web (WWW)*, pages 851–860. ACM.

I. Stewart. 2014. Now We Stronger Than Ever: African-American Syntax in Twitter. In *Proceedings of the EACL Student Research Workshop*, pages 31–37.

Sugumaran and Voss. 2012. Real-time spatio-temporal analysis of West Nile Virus using Twitter Data. In *Proceedings of Computing for Geospatial Research and Applications*.

Z. Tufekci. 2014. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media*.

Xu, Bhargava, Nowak, and Zhu. 2012. Socioscope: Spatio-temporal signal recovery from social media. In *Proceedings of the European Conference on Machine Learning: PKDD*.