

Profiling Reputation of Corporate Entities in Semantic Space

Notebook for RepLab at CLEF 2012

Jussi Karlgren, Magnus Sahlgren, Fredrik Olsson, Fredrik Espinoza, and Ola Hamfors

Gavagai AB
Skånegatan 97, 116 35 Stockholm
www.gavagai.se
jussi@gavagai.se

Abstract Gavagai used its first-generation baseline system for the profiling task for evaluation campaign for online reputation management systems of CLEF 2012. The system builds on large scale analysis of streaming text and performed excellently on this task with standard settings.

1 Profiling corporate reputation in streaming online data

The profiling task was defined to be based on real data, using a set of microblog posts from Twitter filtered to contain a company name. The experimental data consisted of thirty-six sets of microblog post references, each set potentially relevant to a named company. The first six sets were used for training and the thirty latter ones for testing.

The task was for each post in the test set first to determine whether it refers to the company name mentioned in it (some names are ambiguous and sometimes the company name is mentioned in passing) and then assess whether the tweet improves or detracts from the reputation of the company. This task is close to but not identical to a typical negative-positive sentiment analysis. Firstly, the tweets may not be attitudinal but factual, yet retain implications for a company's reputation: a factual report on the state of the world may impinge negatively or positively on a company name. Secondly a statement of attitude couched in ever so attitudinally loaded terms might not have the effect with respect to polarity on company reputation: a glowingly positive report on the wrong aspect of a product or service might not be what the company wants to represent itself with.

The details of the task are given in the introduction to the Evaluating Online Reputation Management Systems Lab of the 2012 CLEF conference [1].

2 System description

Gavagai provides through its Ethersource suite of services tools for monitoring *targets of interest* for some commercial purpose in streaming data of any scale and editorial quality in any language with respect to *semantic poles* of some permanence. Ethersource is based on distributional semantics [6] represented in a semantic space [5], and realised

through a proprietary implementation of the Random Indexing processing framework [2] as described in our position paper at the recent Online Reputation Management workshop [4]. Ethersource is under constant development and the results from this evaluation are being fed back into the system quality assurance cycle.

A target in Ethersource is defined through manual entry of a number of representative terms. In this case the targets were defined through their primary name (“lufthansa”, “#lufthansa”, “lufthansa’s”) and a small number of additional or blocking terms obtained through a support system based on a semantic space model built from previous large scale analysis of streaming text (“blackberry” → “#rim”).

A semantic pole in Ethersource is likewise defined through a larger and more permanently selected number of terms. This term set can be extensive or limited, depending on if recall or precision is crucial for the task at hand and if typical expression of this pole is wide-ranging or more exact [7]. For typical sentiment analysis purposes, the poles can be defined through a list of positive and negative terms; for other purposes other word lists can be used — in our commercial context we have a large number of poles and do not generalise to simple positive or negative [3]. For this task, we utilised Gavagai’s standard poles for *customer satisfaction* for English and Spanish, each of some few hundred editorially selected terms, semi-automatically augmented through the semantic space model built from previous large scale analysis of streaming text and static textual collections in each language.

The system took each candidate microblog post as if it were harvested from a live feed, ran it through a standard language identifier, and filtered it through the entity target representation. If the target identifier fired, the post was polarised with respect to the two opposing customer satisfaction poles defined for the language as identified. The polarisation score, normally aggregated by our system over streaming data into a time series and monitored by our customers for change, varies between 0 and 1 and is not designed to make decisions for text items in isolation from their context. In this case, if the score of either pole was over a editorially set threshold the post was considered to have effect on company reputation. If both scores were high, the larger score was used.

3 Results

Our baseline system as it stands achieved excellent results in the evaluation, with an overall profiling accuracy of 40.0%, calculated by both relevance and polarity identification being correct.

In the relevance assessment or filtering process where our system uses simple lexical term recognition as an indicator of relevance the settings we chose yielded the highest reliability score of any system and near highest sensitivity achieving a filtering accuracy of 77.4%. For predicting effect and direction of effect on reputation our system found that the threshold for taking a post into account was set somewhat conservatively yielding an accuracy of 37%.

4 Analysis

This sort of analysis is by necessity very subjective. A post such as

definitivo lo que me choca de armani es su estilismo

might arguably be interpreted to be neutral, negative, or positive for Armani and a post such as

shit bitch me and @name livin in up at tha Marriott BITCH

as negative or positive (hardly neutral) for Marriott depending on one's perspective; a post such as

BA set to by BMI from Lufthansa

can be interpreted to be neutral, negative, or positive for Lufthansa depending on what one knows about BMI. This is the reason we in commercial application of our system work on time series and sequences rather than single posts.

References

1. Amigó, E., Corujo, A., Gonzalo, J., Meij, E., Rijke, M.d.: Overview of RepLab 2012: Evaluating online reputation management systems. In: CLEF 2012 Labs and Workshop Notebook Papers (2012)
2. Kanerva, P.: Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive Computation* 1(2), 139–159 (2009)
3. Karlgren, J., Sahlgren, M., Olsson, F., Espinoza, F., Hamfors, O.: Usefulness of sentiment analysis. In: ECIR 2012, 34th European Conference on Information Retrieval. Barcelona (2012)
4. Olsson, F., Karlgren, J., Sahlgren, M., Espinoza, F., Hamfors, O.: Technical requirements for knowledge representation for attitude mining on a realistic scale. In: Proceedings of the Workshop on Reputation Management in Social Media at LREC'12. Istanbul (2012)
5. Sahlgren, M.: The Word-Space Model: using distributional analysis to represent syntagmatic and paradigmatic relations between words in high-dimensional vector spaces. Ph.D. thesis, Stockholm University (2006), <http://soda.swedish-ict.se/437/>
6. Sahlgren, M.: The distributional hypothesis. *Rivista di Linguistica (Italian Journal of Linguistics)* 20(1), 33–53 (2008), <http://soda.swedish-ict.se/3941/>
7. Sahlgren, M., Karlgren, J., Eriksson, G.: SICS: Valence annotation based on seeds in word space. In: Fourth International Workshop on Semantic Evaluations (SemEval-2007). Association for Computational Linguistics (2007), <http://soda.swedish-ict.se/2593/>