Information Processing and Management 52 (2016) 193–216

Contents lists available at ScienceDirect

Information Processing and Management

journal homepage: www.elsevier.com/locate/ipm

Estimating Reputation Polarity on Microblog Posts

Maria-Hendrike Peetz[a,*], Maarten de Rijke[a], Rianne Kaptein[b]

[a] University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
[b] TNO, Brassersplein 2, 2612 CT Delft, The Netherlands

ARTICLE INFO

Article history:
Received 6 August 2013
Revised 16 May 2015
Accepted 7 July 2015
Available online 6 November 2015

Keywords:
Social media analysis
Online reputation analysis

ABSTRACT

In reputation management, knowing what impact a tweet has on the reputation of a brand or company is crucial. The reputation polarity of a tweet is a measure of how the tweet influences the reputation of a brand or company. We consider the task of automatically determining the reputation polarity of a tweet. For this classification task, we propose a feature-based model based on three dimensions: the source of the tweet, the contents of the tweet and the reception of the tweet, i.e., how the tweet is being perceived. For evaluation purposes, we make use of the RepLab 2012 and 2013 datasets. We study and contrast three training scenarios. The first is independent of the entity whose reputation is being managed, the second depends on the entity at stake, but has over 90% fewer training samples per model, on average. The third is dependent on the domain of the entities. We find that reputation polarity is different from sentiment and that having less but entity-dependent training data is significantly more effective for predicting the reputation polarity of a tweet than an entity-independent training scenario. Features related to the reception of a tweet perform significantly better than most other features.

1. Introduction

Social media monitoring and analysis has become an integral part of the marketing strategy of businesses all over the world (Mangold & Faulds, 2009). Companies can no longer afford to ignore what is happening online and what people are saying about their brands, their products and their customer service. With growing volumes of online data it is infeasible to manually process everything written online about a company. Twitter is one of the largest and most important sources of social media data (Jansen, Zhang, Sobel, & Chowdury, 2009). Tweets can go viral, i.e., get retweeted by many other Twitter users, reaching many thousands of people within a few hours. It is vital, therefore, to automatically identify tweets that can damage the reputation of a company from the possibly large stream of tweets mentioning the company.

Tasks often considered in the context of online reputation management are *monitoring* an incoming stream of social media messages and *profiling* social media messages according to their impact on a brand or company's reputation. We focus on the latter task. In particular, we focus on the problem of determining the *reputation polarity* of a tweet, where we consider three possible outcomes: positive, negative, or neutral. Knowing the reputation polarity of a single tweet, one can either aggregate this knowledge to understand the overall reputation of a company or zoom in on tweets that are dangerous for the reputation of a company. Those tweets need counteraction (van Riel & Fombrun, 2007).

* Corresponding author.
E-mail addresses: m.h.peetz@uva.nl (M.-H. Peetz), derijke@uva.nl (M. de Rijke), rianne.kaptein@tno.nl (R. Kaptein).

http://dx.doi.org/10.1016/j.ipm.2015.07.003
0306-4573/

The reputation polarity task is a classification task that is similar to, but different in interesting ways, from *sentiment analysis*. For example, a post may have a neutral sentiment but may be negative for reputation polarity. Consider, for instance, the statement *The room wifi doesn't work.*, which is a factual statement that may negatively impact the reputation of a hotel.

There are two standard benchmarking datasets for reputation polarity, the RepLab 2012 dataset (Amigó, Corujo, Gonzalo, Meij, & de Rijke, 2012a) and the RepLab 2013 dataset (Amigó et al., 2013), made available as part of RepLab, a community-based benchmarking activity for reputation analysis. In view of the distinction that we have just made between sentiment analysis and reputation polarity, it is interesting to observe that the best performing reputation polarity classifiers at RepLab are sentiment-based. The main research question we address is:

RQ1 Can we improve the effectiveness of baseline sentiment classifiers by adding additional information?

The RepLab 2012 and 2013 datasets have different training and testing scenarios: the 2012 dataset uses a training and testing setup that is independent of individual brands or companies ("entities"), while this dependence is introduced in the 2013 dataset. We ask:

RQ2 How do different groups of features perform when trained on entity-(in) dependent or domain-dependent training sets?

Our last research question is exploratory in nature. Having introduced new features and interesting groups of features, we ask:

RQ3 What is the added value of features in terms of effectiveness?

Without further refinements, RQ3 is a very general research question. One of the contributions of this paper, however, is the way in which we model the task of determining the reputation polarity of a tweet as a three-class classification problem: we build on communication theory to propose three groups of features, based on the *sender* of the tweet, on the *message* (i.e., the tweet itself), and on the *reception* of the message, that is, how the tweet is being perceived.

While we use and compare some features that are known from the literature (Naveed, Gottron, Kunegis, & Alhadi, 2011), a second contribution that we make in this paper consists of new features to capture the reception of messages-this is where the difference between reputation polarity and sentiment analysis really shows.

Furthermore, as we will see below, reputation polarity class labels are highly skewed and data for some features is missing; our third contribution below consists of an analysis of sampling methods to alleviate the problem of skewness.

Another important contribution that we make concerns the way in which we operationalize the reputation management task. Social media analysts use company-specific knowledge to determine the reputation polarity (Corujo, 2012). In line with this, we discover that sets of tweets pertaining to different entities may be very different in the sense that different features are effective for modeling the reputation polarity. We therefore provide an operationalization of the reputation polarity task using the RepLab 2012 dataset in which we train and test on company-dependent datasets instead of using a generic training set. We find that we can avoid overtraining and that training on far fewer data points (94.4% less) per entity gives up to 37% higher scores. The observation transfers to the RepLab 2013 dataset which is operationalized in precisely that way.

Finally, this paper adds a new point of view for the business analysis perspective: here our biggest contribution is the difference in performance of features when trained on entity or domain dependent or independent data. Features pertaining to the author of the message seem to be generalizable while others are not.

We proceed with a definition of the reputation polarity task in Section 2. Section 3 introduces our features and reputation polarity model. We detail our experiments, results and analysis in Sections 4 and 5, respectively. Section 6 provides an overview of related work, and we conclude in Section 7.

2. Task definition

The current practice in the communication consultancy industry is that social media analysts manually perform labeling and classification of the content being analyzed (Amigó et al., 2012a). Two of the most labor intensive tasks for reputation analysts are *monitoring* and *profiling* of media for a given company, product, celebrity or brand ("entity"). The monitoring task is the (continuous) task of observing and tracking the social media space of an entity for different topics and their importance for the reputation of the entity. Here, the retrieval and aggregation of information concerning the entity is most important. Technically, the monitoring task can be understood as consisting of two steps as follows:

(*Cluster*) cluster the most recent social media posts about an entity thematically, and
(*Rank*) assign relative priorities to the clusters.

In this paper we focus on the profiling task, which is the (periodic) task of reporting on the status of an entity's reputation as reflected in social media. To perform this task, social media analysts need to assess the relevance of a social media