Decision Support Systems 88 (2016) 67–75

Contents lists available at ScienceDirect

Decision Support Systems

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

Negation scope detection in sentiment analysis: Decision support for news-driven trading

Nicolas Pröllochs*, Stefan Feuerriegel, Dirk Neumann

Chair for Information Systems Research, University of Freiburg, Platz der Alten Synagoge, 79098 Freiburg, Germany

ARTICLE INFO

Article history:
Received 13 November 2015
Received in revised form 28 May 2016
Accepted 28 May 2016
Available online 5 June 2016

Keywords:
Decision support
Machine learning
Sentiment analysis
Negation scope detection
Financial news

ABSTRACT

Decision support for financial news using natural language processing requires robust methods that process all sentences correctly, including those that are negated. To predict the corresponding negation scope, related literature commonly utilizes rule-based algorithms and generative probabilistic models. In contrast, we propose the use of a tailored reinforcement learning method, since it can conquer learning task of arbitrary length. We then perform a thorough comparison with a two-pronged evaluation. First, we compare the predictive performance using a manually-labeled dataset. Here, reinforcement learning outperforms common approaches from the related literature, leading to a balanced classification accuracy of up to 70.17%. Second, we examine how detecting negation scopes can improve the accuracy of sentiment analysis for financial news, leading to an improvement of up to 10.63% in the correlation between news sentiment and stock market returns. This reveals negation scope detection as a crucial leverage in decision support from sentiment.

1. Introduction

Sentiment analysis is a frequently utilized approach for sensing the tone in written content. Although this method is popular in various domains, it has gained traction especially in the financial domain. Here, it can relate positive and negative wording in financial disclosures to subsequent stock market movements [1–4]. While this idea seems intriguing, previous research identifies obstacles: *"positive tone is difficult to accurately measure, since positive words are easily negated in ways difficult to programmatically identify"* [5]. In fact, negations are a fundamental stylistic device for inverting the meaning of both words and sentences. For instance, financial news in the form of German ad hoc announcements contains negations in as many as 4.74% of all sentences.

As such, identifying and predicting negated parts are crucial for providing accurate decision support, since otherwise sentences are likely to be classified erroneously. According to Ref. [5], *"the results for positive words are mixed because many times, negative phrases are wrapped in positive words"*. To alleviate such a possible cause of disturbance, this paper proposes and evaluates algorithms with the aim of predicting negated parts of sentences. This demonstrates how the narrative content of disclosures [6–10] can be harnessed to provide *decision support* for investors.

Negations can appear in various forms; inverting not only the meaning of single words but also of whole phrases. Accordingly, one refers to the part whose meaning is changed as the so-called *negation scope*. Furthermore, negations can also flip the meaning of sentences implicitly; e.g. *"the company has invented a new product; it was the first and last time"*. Unfortunately, many machine learning algorithms struggle with this type of problem as it is virtually impossible to encode with a fixed-length vector while preserving its order and context. To justify this formally, we immediately observe that both negation scopes and sentences can be of changing length m that can theoretically range from one to infinity. Each word w_i, $i \in \{1, \ldots, \infty\}$, thus represents an individual classification task

$$f : (w_i, [w_1, w_2, \ldots]) \mapsto \{\text{Negated}, \neg\text{Negated}\}, \tag{1}$$

where $[w_1, w_2, \ldots]$ is an ordered list of arbitrary length providing context information. This learning task f lies in contrast to many algorithms from classical machine learning, such as supervised learning, which always require an input vector of a fixed, pre-defined length m. As a remedy, this calls for special methods in order to deal with natural text and negation scope detection in particular.

Previous approaches for negation scope detection can be primarily grouped into two categories (e.g. [11]). On the one hand,

* Corresponding author. Tel.: +49 761 203 2395; fax: +49 761 203 2416.
E-mail addresses: nicolas.proellochs@is.uni-freiburg.de (N. Pröllochs), stefan.feuerriegel@is.uni-freiburg.de (S. Feuerriegel), dirk.neumann@is.uni-freiburg.de (D. Neumann).

http://dx.doi.org/10.1016/j.dss.2016.05.009
0167-9236/

rule-based approaches are straightforward to implement. However, they can barely identify implicit negations, not to mention linguistic peculiarities, such as sarcasm or irony. On the other hand, generative probabilistic models are computationally more expensive and require ex ante transition probabilities but can adapt to domain-specific features/language. Examples of the latter include Hidden Markov models or conditional random fields.

As an alternative, we propose the use of a third group, namely, *reinforcement learning*. The reason is that this method can work well with learning tasks of arbitrary length [12]. Interestingly, reinforcement learning has been greatly neglected in the domain of natural language processing (see Section 2), even though it aims at replicating actual human negation detection. More precisely, it learns based on past experience using only limited feedback in the form of a reward signal that indicates how well the learner is behaving while not explicitly detailing how to improve its behavior. In addition, the episode-based and flexible structure of this method can handle highly complex sentences and thus appears to be well suited for negation scope detection.

As a main contribution, this paper explores the impact of negation scope detection on the sentiment analysis of financial news. This allows us to investigate and compare predicted negation scopes in practice. In contrast to previous research, we not only implement a single method but also provide a holistic comparison of different options: rule-based approaches and Hidden Markov models, as well as conditional random fields, with both supervised and unsupervised learning. In addition, we propose the use of reinforcement learning in order to predict negation scopes. Moreover, we also investigate the sensitivity to the underlying sets of negation words. Finally, all approaches are compared in terms of their predictive performance and computational time using two dimensions: (1) we examine the accuracy of a manually-labeled dataset consisting of negated sentences from financial news. (2) Since the chosen labels are biased by subjective interpretation, the stock market reaction of investors serves as an objective measure. Here, we compute sentiment values for financial news and then compare the correlation of these sentiment values with the corresponding stock market returns. Altogether, this two-pronged evaluation reveals a compelling way of measuring and improving the accuracy of sentiment analysis.

The rest of this paper is structured as follows. Section 2 provides an overview of related works on negation scope detection, while Section 3 explains how we improve the state of the art with reinforcement learning. We then benchmark these approaches with financial news in Section 4. Finally, Section 5 discusses our results and identifies the implications of our research.

2. Related work

This section recapitulates related research in a structured fashion and points out how sentiment analysis can be improved by accurate negation detection. The detection of negation scopes essentially affects any context [13–17] in which sentiment analysis occurs but is particularly true when it comes to financial news [1, 5, 18]. Managers can easily negate words in ways that are difficult to identify computationally [1, 5]. Investigations about 10-K reports indicate that companies often frame negative news using positive words and rarely relay positive news in terms of negative words [18]. A thorough survey on the role of negation in sentiment analysis, as well as how to embed negation scopes, is found in Ref. [14].

Negation scope detection based on heuristic rules results in significant improvements when working with user-generated blog entries [19]. A similar approach inverts the meaning only within windows of a fixed size [20], which is evaluated using IMDb movie reviews. Investigations of political news during an election campaign incorporate part-of-speech tags to capture entity interactions by modeling subjects and objects [16]. Efforts in opinion mining systems on several manually-labeled corpora lead to robust cross-domain performance of rule-based approaches [21]. Research using similar approaches on the same data demonstrates problems with implicit negations and recommends the application of machine learning approaches [22].

In the domain of machine learning, authors propose the use of conditional random fields [13] to predict negation scopes for the sentiment analysis of product reviews. These are, in fact, a common [17, 23, 24, 25] methodological choice. Even though Ref. [25] compares two approaches, namely, conditional random fields and an approach based on regular expressions, the author does not explicitly consider financial news nor sets of negation words. In Ref. [26], the authors hand-checked a small set of negation scopes, arguing that the denial component is irrelevant when studying filings from initial public offerings (IPO). Finally, we note that negation scope detection is also a frequent research topic (e.g. [11]) for information retrieval aimed at medical reports.

Hence, this paper addresses the following research goal: we propose and compare algorithms to predict negation scopes by utilizing both linguistic rules and machine learning from previous work. In addition, we also evaluate the suitability of reinforcement learning for predicting negation scopes. Although this method is particularly popular in the domain of robotics and game theory (e.g. [27]), it is only rarely applied to the domain of computational linguistics. In fact, we are not aware of any publication that utilizes reinforcement learning to predict negation scopes. Moreover, to our best knowledge, this is the first study that predicts negation scopes designed to enhance the sentiment analysis of financial news.

3. Methods and materials

This section introduces our research methodology, as well as our datasets. Fig. 1 depicts how, in a first step, each announcement is subject to preprocessing steps which transform the running text into machine-readable tokens. Among others, this includes standard techniques, such as tokenization, part-of-speech tagging and stemming [28]. Afterwards, we extend previous works [29] and evaluate the suitability of the reinforcement learning approach for negation scope detection with common methods from the related literature. This includes generative probabilistic models, such as Hidden Markov models and conditional random fields, as well as rule-based algorithms.

We follow a two-sided approach. First, we evaluate the predictive performance on a manually-labeled dataset. Second, we examine the benefits of predicted negation scopes in an application from natural language processing; more precisely, how the sentiment analysis of financial news can be improved. In order to calculate the polarity of news announcements, we utilize the frequently employed approach of the so-called *Net-Optimism* sentiment [3]. This approach assumes that a news announcement with a positive message correlates with more positive words and vice versa. Thus, the Net-Optimism measure rates the content according to the frequencies of words classified in pre-defined dictionaries as either *positive* or *negative*. Although many other approaches can be found in the literature [30, 31], Net-Optimism provides not only robust results [32], but its simplicity makes later comparisons straightforward. We combine this measure with negation scope detection by inverting the polarity of words that are negated. For example, in the sentence *"company sales have ⟨not significantly increased⟩"*, the negation scope is given by the sentence fragment inside the angled brackets. Here, the meaning of *significantly increased* is negated by the negation word *not* and the connotation according to the pre-defined dictionaries should be reversed. In the following sections, we present the underlying methods in order to determine negation scopes.