

A COMPARATIVE STUDY OF THE RELEVANCE OF INDIRECT AND DIRECT OPINIONS IN ECONOMIC TEXTS

MUSAT, C[laudiu] C[ristian] & TRAUSAN - MATU, S[tefan]

Abstract: We propose a new approach to opinion mining that relies not on explicit occurrences of opinions in economic texts but on implicit ones extracted from the polarity of the discussed economic subtopic. We prove that the majority of economic texts contain personal opinions that can be retrieved based solely on the implicit opinions within. We rely on co-occurrences of economic indicators and future state modifiers to determine the polarity of each text. We benchmark the method's results with those obtained using SentiWordNet.

Key words: opinion mining, economics, SentiWordNet

1. INTRODUCTION

Generally the focus in opinion mining has been on analyzing occurrences of direct opinions in the studied texts. However there are domains where this approach is less adequate, where we need to broaden this perspective in order to detect opinions that only derive from context and are not expressed directly. We limit the analysis to economic documents. The paper continues by discussing the state of the art, the proposed system is described in Section 3, while results and conclusions follow in Sections 4 and 5.

2. STATE OF THE ART

One of the most influential works in the field was that of Pang and Lee (2008), where documents were classified according to the overall opinions expressed. Wiebe et al. (2005) discriminated between subjective and objective sentences. To achieve better granularity the focus shifted to determining the polarity of expressed opinions, ranging from a positive/negative dichotomy to more complex classifications based on opinion strength (Chen and Chiu, 2009). SentiWordNet – SWN - (Esuli and Sebastiani, 2006) was created to assign a polarity to all words in the WordNet (Fellbaum, 1998) dictionary. All the methods above rely on the presence of explicit opinions in the analyzed texts

Regarding economics, in a study regarding sentiment analysis on financial blogs, Ferguson et al. (Ferguson et al. 2009) showed that using paragraph level annotations and a Bayesian classifier, more than 60% of the analyzed texts contained personal opinions. However their study was restrained to texts that explicitly referred to companies listed in the S&P 500 index. Mainstream media has been addressed by Ahmad et al. (2006). Furthermore opinions expressed in the press have been linked to relations between listed companies and to their stock prices (Ku et al. 2006). The lack of explicit opinions does not always imply the lack of an overall opinion. Amigo et al. (2009) showed that the mere mention of a product in a review can express positive sentiment towards it.

3. PROPOSED SYSTEM

We propose a method to extract opinions expressed in economic texts that are presented as facts. Economic

predictions were analyzed as certain combinations of economic indicators such as “unemployment”, “productivity” or “markets” and phrases (modifiers) that suggest a certain future evolution. The indicators are divided into two subsets. One that contains indices that rise along with the economy as a whole and one with indices which decline during a period of economic boom. We will call the first “positive” indicators, and the latter “negative”. Modifiers are a collection of n-grams that indicate growth or decrease of the indicator they are referring to. As with the indicators above, the set of modifiers will be split into positive and negative.

The economic terms and phrases were obtained by expanding the EconomyWatch glossary, and the modifiers by expanding an initial handpicked set of words signaling an increase or decrease. The first expansion relies on WordNet (Fellbaum, 1998) synsets and WordNet terms whose glosses contain occurrences of terms in the initial set. The second extraction method employs a Contextual Network Graph (Cegłowski et al. 2003) that retrieves terms that tend to appear together with ones in the initial set in context.

In this experiment we used the publicly available articles from The Telegraph Finance section and its subsections. A total of 21106 articles published between 2007 and 2010 were processed and within the corpus we have found 499 articles that were labeled as having a negative bias and 216 as having a positive one.

The following task is to find the polarity of the author's economic forecast from co-occurrences of economic indicators and future state modifiers. A good example is that of recent discussions about the rise of unemployment. The term unemployment is itself a negative indicator. When the modifier attached to it is a positive one, such as “will grow” the overall prediction is negative. Likewise, if the modifier is negative, like “fell” the result is positive. A simple formula summarizes the above:

$$P(O) = P(I) * P(M) \quad (1)$$

where

- P is the polarity of the construct
- O the resulting opinion
- I is the indicator involved
- And M the attached modifier

Not all pairs encountered are expressions of implicit opinions. The phrase “rising unemployment could prove harmful for the local economy” does not signal a prediction about the evolution of the indicator. While bearing in mind that oversimplification could significantly reduce the opinion extraction system's recall, we used two heuristics to avoid uncertainties similar to the mentioned one.

The first imposed limitation is that the indicator should be either a noun (or, if the indicator is not a unigram, an n-gram containing a noun) or an adjective attached to a neutral noun. Neutral associate nouns are very frequent in the encountered

economic texts, forming constructs such as “economic data”, “unemployment figures” or “mortality numbers”.

The modifier is limited to the associated verb or, when the verb suggests a continuous state of fact, an adverb following that verb, as in “the recovery remains sluggish”. For simplicity we will refer to the entire family of verbs related to the one in the previous example as *continuers*. In order to apply the heuristic adjustments above, a part-of-speech tagger was used.

Term	POS	Requirements
Indicator	Noun/Noun Phrase	None
Indicator	Adjective	Neutral Noun
Modifier	Verb	None
Modifier	Adverb	Continuer Verb

Tab. 1. Relevant occurrences of key terms

Also, detecting negations of the selected indicators and modifiers is crucial for the accuracy of the system as the use of negated terms is often more prevalent than their use without negation.

A document in which the majority of pairs indicate the growth of a positive indicator or the decrease of a negative one will be labeled as having a positive polarity, whereas a document in which the majority of pairs indicate the growth of a negative indicator or the decrease of a positive one will have a negative polarity.

4. RESULTS

The results of the opinion extraction test phase and the precision and recall of the experiment are presented in Tables 2 and 3 next to the SentiWordNet results on the same corpus. Note that there are articles in the root sections of both the positive and negative batches, so that the sum of the articles in the subsections does not equal the total. Both the positive and negative sections contain three subsections: jobs, markets and comments. We denote the positive outcomes as “+”, the negative ones as “-“ and the neutrals as “=”.

Corpus	All	+	-	=	SWN +	SWN -	SWN =
Neg.							
All	499	86	252	161	207	288	4
Jobs	54	1	35	18	19	35	0
Markets	253	44	128	81	100	152	1
Com.	153	32	73	48	60	91	2
Pos.							
All	216	87	50	79	63	94	59
Jobs	14	4	3	7	4	7	3
Markets	152	71	31	50	42	69	41
Com.	34	11	9	14	11	15	8

Tab. 2. Polarity identification test results

5. CONCLUSIONS

The SWN recall is consistently close to 100% because the polarity of all present words is considered and the sum is rarely close enough to 0 to be considered “neutral”. However the precision is consistently smaller than the indirect opinion based system. Also, our system’s recall is linked to the size of the training corpus, the worst results being obtained for the Jobs and Comment sections of the positive batch, where the number of articles is the lowest. Also we note that the comment sections for both the positive and negative articles present a lower accuracy compared to the Markets and Jobs sections because of

the much larger number of subjects discussed, from personal finance to macroeconomics.

Corpus	Precision	Recall	SWN Precision	SWN Recall
Total	71.37	66.43	53.83	91.19
Negative				
All	74.56	67.74	58.18	99.20
Jobs	97.22	66.67	64.81	100.00
Markets	74.42	67.98	60.32	99.60
Comments	69.52	68.63	60.26	98.69
Positive				
All	63.50	63.43	40.13	72.69
Jobs	57.14	50.00	36.36	78.57
Markets	69.61	67.11	37.84	73.03
Comments	55.00	58.82	42.31	76.47

Tab. 3. Polarity identification test precision and recall

By switching focus from explicit opinions in economics to the quantification of implicit predictions of future economic outcomes we hope to have created a novel tool for opinion mining, complementary to existing ones. Future extensions include the processing of irony and metaphors, both widely used in financial articles have not been addressed in the current experiment. In addition to that, there are numerous cases where outside quotes are used, only to be disagreed with by the author. This too is a type of negation that needs to be included in the opinion extraction module. But perhaps The most important of all future improvements is the addition of more indicators and modifiers to improve the system’s recall.

6. ACKNOWLEDGEMENTS

Supported by the EU Research Grant POSDRU 7713.

7. REFERENCES

- Ahmad K.; Cheng D. & Almas Y. (2006). Multi-lingual sentiment analysis of financial news streams. First International Workshop on Grid Technology for Financial Modeling and Simulation: pp. 984-991
- Amigo E.; Spina D. & Bernardino B. (2009). User Generated Content Monitoring System Evaluation. Proceedings of WOMSA 2009. pp 1-13
- Ceglowski M. & Coburn A. (2003). Semantic Search of Unstructured Data using Contextual Network Graphs.
- Chen L.S. & Chiu H.J. (2009). Developing a Neural Network based Index For Sentiment Classification. Proceedings of IAENG 2009: pp 744-749.
- Cruz F.; Troyano J.; Ortega J. & Enriquez F. (2009). Domain Oriented Opinion Extraction Methodology. Proceedings of WOMSA. pp. 52-62
- Esuli, A.; Sebastiani F. (2006). SENTIWORDNET: A Publicly Available Lexical Resource for Opinion Mining, Proceedings of LREC, pp. 417-442, Genoa, Italy, may 2006
- Felbaum C. (1998). WordNet: An Electronic Lexical Database. Cambridge, MA: MIT Press
- Ferguson P.; O’Hare N.; Berminghamv A.; Tattersall S.; Sheridan.; Gurrin C. & Smeaton A. (2009). Exploring the use of Paragraph-level Annotations for Sentiment Analysis of Financial Blogs. Proceedings of WOMSA 2009. 42-52
- Liu B. (2008). Opinion Mining. In proceedings of WWW-2008
- Pang B. & Lee L. (2008). Opinion Mining and Sentiment Analysis, Foundation and Trends in Information Retrieval, 2: 1-135
- Wiebe J.; Wilson T. & Cardie C. (2005). Annotating Expressions of Opinions and Emotions in Language. Language Resources and Evaluation, 39(2):165-210