

Rule-Based Sentiment Analysis in Narrow Domain: Detecting Sentiment in Daily Horoscopes Using *Sentiscope*

Željko AGIĆ¹ Danijela MERKLER²

(1) Department of Information and Communication Sciences

(2) Department of Linguistics

Faculty of Humanities and Social Sciences, University of Zagreb

Ivana Lučića 3, 10000 Zagreb, Croatia

{zagic,dmerkler}@ffzg.hr

ABSTRACT

We present a prototype system — named *Sentiscope* — for collecting daily horoscopes from online news portals written in Croatian, detecting polarity phrases and overall sentiment conveyed by these texts and providing sentiment-analysis-based visualizations in a graphical user interface on the web. The system was evaluated using a dataset of daily horoscopes which was manually annotated for (positive and negative) polarity phrases and (positive, negative and neutral) overall sentiment. Linearly weighted kappa coefficient of 0.593 has indicated moderate inter-annotator agreement on overall sentiment annotation. The system achieved an F_1 -score of 0.566 on overall sentiment and 0.402 on phrase detection. An overview of implementation is provided — with special emphasis on the polarity phrase detection module implemented in NooJ linguistic IDE — and the system is made available to users on the web.

TITLE AND ABSTRACT IN CROATIAN

Analiza sentimenata pravilima u uskoj domeni: pronalaženje sentimenata u dnevnom horoskopu sustavom *Sentiscope*

Predstavljamo prototip sustava — nazvanoga *Sentiscope* — za prikupljanje dnevnih horoskopa s novinskih internetskih portala pisanih hrvatskim jezikom, pronalaženje polarnih izraza i ukupnih sentimenata prenesenih tim tekstovima i pružanje skupa vizualizacija zasnovanih na analizi sentimenata putem internetskoga grafičkog korisničkog sučelja. Sustav je vrjednovan s pomoću skupa dnevnih horoskopa u kojima su ručno označeni (pozitivni i negativni) polarni izrazi i (pozitivni, negativni i neutralni) ukupni sentiment. Linearni je *kappa*-koeficijent od 0.593 ukazao na umjereno slaganje označitelja pri označavanju ukupnoga sentimenta. Točnost je sustava izražena F_1 -mjerom od 0.566 pri pronalaženju ukupnoga sentimenta i 0.402 pri pronalaženju polarnih izraza. Dan je pregled izvedbe sustava — s posebnim naglaskom na modulu za pronalaženje polarnih izraza izrađenom s pomoću lingvističkoga razvojnog okruženja NooJ — i korisnicima je omogućen internetski pristup sustavu.

KEYWORDS: sentiment analysis, narrow domain, rule-based system.

KEYWORDS IN CROATIAN: analiza sentimenata, uska domena, sustav temeljen na pravilima.

1 Introduction and related work

Sentiscope is a prototype system for sentiment analysis in daily horoscopes written in Croatian. It crawls the Croatian web on a daily basis and collects horoscope texts from several specialized websites and daily news portals. The texts are processed with a manually designed rule-based module for polarity phrase detection. The texts are then assigned with overall sentiment scores which are calculated by counting polarity phrases. The results of semantic processing are stored and the texts with the respective annotations of both polarity phrases and the overall sentiments are provided to users via a graphical user interface in the form of a web application.

Implementation of *Sentiscope* draws from the work on approaches to sentiment analysis in financial texts and related work on sentiment analysis presented in, e.g., (Ahmad et al., 2005, 2006a,b; Almas and Ahmad, 2007; Devitt and Ahmad, 2007, 2008; Daly et al., 2009; Remus et al., 2009). More specifically, drawing from the experiment with rule-based sentiment analysis in financial reports written in Croatian presented in (Agić et al., 2010) — which resulted with a high precision prototype system — and the previously mentioned work on rule-based sentiment analysis in general, we attempted to approach the problem of sentiment analysis in Croatian text from a very specific, narrow and expectedly difficultly processable genre, i.e., horoscope text from the web.

Alongside system implementation and evaluation, we emphasize the ambiguity of sentiment detection in general — end especially in narrow and ambiguous domains, represented here by horoscope text — by creating a manually annotated dataset of horoscopes and calculating inter-annotator agreement for the overall article sentiment manual annotation task. This special emphasis is motivated by previous explorations of properties of various sentiment analysis challenges, relating inter-annotator agreement and task difficulty, such as (Pang and Lee, 2008) and (Bruce and Wiebe, 1999; Wiebe et al., 1999, 2004; Shanahan et al., 2006). For example, it is specifically stated by (Pang and Lee, 2008) that "different researchers express different opinions about whether distinguishing between subjective and objective language is difficult for humans in the general case." They also state that "for example, (Kim and Hovy, 2006) note that human annotators often disagreed on whether a belief statement was or was not an opinion while other researchers have found inter-annotator agreement rates in various types of subjectivity classification tasks to be satisfactory." Here we implicitly address the relation between difficulty of manual sentiment annotation and meaningfulness of tackling the same annotation problem algorithmically. Moreover, following (Riloff et al., 2003; Wiebe, 2000; Wilson et al., 2005), we investigate the role of certain parts of speech — such as adjectives, adverbs, nouns and verbs — in detecting different classes of polarity phrases.

In the following sections, we describe the system implementation and evaluation on the tasks of detecting polarity phrases and detecting overall article sentiment. The system prototype is available on the web (<http://lt.ffzg.hr/sentiscope/>).

2 System implementation

System overview is given in Figure 1 (left side). The system is basically a web- and Linux-based application built by open source technologies and it consists of four main components:

1. the focused web crawler written in PHP that collects and stores horoscopes from a number of Croatian horoscope and daily news portals,
2. the rule-based sentiment detector that detects positive and negative polarity phrases




Figure 1: System overview and main polarity phrase detection grammar

2012-07-18 **ovan** ■■■■■

■ *24sata.hr* Uznemirenost i loši predošjećaji činit će vas **nemirnima** cijeli dan. Intuitivno ćete izbjegavati blizinu osoba za koju osjećate da šire **negativnu** energiju. **Napeti** obiteljski odnosi narušavat će vaš unutarnji mir, što će vas prisiliti na neke **odgadane** odluke. Patite li od alergije, izbjegavajte **štetne** utjecaje iz okoliša.

■ *dalmacijanews.com* Konačno ćete moći postaviti pravu dijagnozu za sve svoje neuspjehe. Bit ćete prisiljeni potisnuti neke svoje **vrline**, za dobrobit sebe i svojih **najdražih**. Konačno ćete shvatiti kako je dosta postavljanja ljubavi na prvo mjesto, kad vam nije uzvraćena. Poželjet ćete za sebe **samo najbolje** i u tome **ćete uspjeti**. Ljubav vas vodi po labirintu. Saznat ćete **zanimljive** događaje o osobi **koju volite**.

Figure 2: Screenshot of the user interface

in horoscope text and is implemented as a set of local grammars designed in the NooJ linguistic development environment (Silberztein, 2004, 2005),

3. overall sentiment detector written in PHP that estimates overall article sentiment, i.e., horoscope sentiment by counting positive and negative polarity phrases and
4. the graphical user interface for assessing sentiment-annotated daily horoscopes and sentiment statistics over periods of time, as illustrated by Figure 2 and 4.

All horoscopes, respective polarity phrase annotations and overall sentiment scores are stored in a MySQL database. The user interface currently provides daily horoscopes with in-line annotations for all twelve zodiac signs (see Figure 2) and historical data in the form of overall sentiment diagrams. Both visualizations also conveniently and entertainingly serve as indicators of sentiment inconsistencies across zodiac signs and web sources. However, regardless of the overall purpose (or purposelessness) of such texts, it is shown here that texts from the specific horoscope genre written in Croatian are very difficult to process with respect to sentiment annotation and thus deserving the given research focus.




Figure 3: Example of positive polarity phrase detection using NooJ local grammars — hr. *nećete biti razočarani* (en. *you will not be disappointed*)

As mentioned previously, overall article sentiment is estimated from the number of detected phrases denoting positive or negative sentiment. Currently, articles are tagged as positive if the number of positive phrases is greater than the number of negative phrases contained within them and vice versa. If their counts are equal, the article is tagged as neutral. Polarity phrase detection is done by using a series of rules in form of local grammars or lexical finite state transducer cascades implemented in NooJ linguistic development environment, as illustrated in Figure 3.

Rules were designed in two stages — first from scratch and then by observing a development set of horoscope texts. For development and testing, we have collected horoscopes from seven largest Croatian websites containing daily horoscopes as indicated by the Google search index. Horoscopes were collected from 2012-02-11 to 2012-05-10. 7,716 articles with 484,179 tokens were collected. 333 articles were chosen for the development set and were manually annotated for overall sentiment and polarity phrases. Observed agreement of 75.97% on overall sentiment annotation was measured between the two annotators. The kappa coefficient indicated good strength of this agreement (0.641), while the linearly weighted kappa coefficient (0.593) assessment indicated moderate agreement. The stats are given in Table 1 and they indicate that the disagreement between the annotators was distributed almost exclusively within the category of neutral sentiment articles. The annotators agreed on positive sentiment in 80.69% of the annotations, while the observed agreement was 82% on negative sentiment and 66.09% on neutral sentiment. If we were to entirely exclude the category of neutral sentiment from data in Table 1, the observed agreement would be 99.44% and the respective kappa coefficient would amount to 0.989 and thus represent very good agreement strength.

Table 2 emphasizes the relation between the polarity phrases detected in articles and the overall sentiment of the articles and as such, it is the theoretical baseline for building a system that estimates overall sentiment of text from the number and type of polarity phrases that it contains. The table shows that the positive sentiment articles tend to contain much more positive polarity phrases, as 71.80% of the positive polarity phrases was found in positive sentiment articles, as opposed to 3.33% in negative and 24.87% in neutral sentiment articles. The same was found to

	+	-	x	Σ
+	94	0	26	120
-	1	82	31	114
x	18	4	77	99
Σ	113	86	134	333

Table 1: Inter-annotator agreement on overall sentiment

	<p>	<n>	both	<p> in both	<n> in both
+	410	27	23	85	27
-	19	321	15	19	53
x	142	145	67	117	115

Table 2: Relation between overall article sentiment (+, -, x) and polarity phrases (<p>, <n>)

apply for negative polarity phrases as well: 65.11% of them were located in negative sentiment articles, 5.48% in positive sentiment articles and 29.41% in articles carrying neutral overall sentiment. This justified a system design in which polarity phrases are counted in articles and overall sentiment assigned from the polarity group with the highest count. In addition to this, Table 2 also shows the number of articles in which both positive and negative polarity phrases were observed (table column *both*), along with separate counts of positive and negative polarity phrases (table columns *<p> in both* and *<n> in both*) for these articles. The distribution further supports the system design, being that positive polarity phrases are once again predominant in positive sentiment articles (75.89% positive vs. 24.11% negative) and negative polarity phrases dominate in negative sentiment articles (73.61% negative vs. 26.39% positive) while they are almost evenly spread in neutral sentiment articles (50.43% positive vs. 49.57% negative).

Rules for polarity phrase detection are grouped in two NooJ local grammars — one for positive sentiment and one for negative sentiment detection (see Figure 1, right side). Each of these grammars consists of lists of words and phrases for three parts-of-speech: adjectives, nouns and verbs. Another part-of-speech generally considered important in sentiment analysis — adverbs — are included within adjectives, due to the specifics in Croatian morphology, i.e., the fact that many adverbs in Croatian are homographic with adjective forms in singular nominative case in neuter gender: e.g., *brzo dijete* (en. *fast child*) *brzo trči* (en. *runs fast*). Words and phrases are manually derived from a number of daily horoscopes and — except for the characteristic key words and key phrases for the horoscope domain — there is a number of domain independent words and phrases, e.g., *dobro* (en. *good*), *izvršno* (en. *great*), *odlično* (en. *excellent*) for positive sentiment, and *loše* (en. *bad*), *slabo* (en. *weak*), *nedovoljno* (en. *unsatisfying*) for negative sentiment. We derived 170 words and phrases for negative and 139 words and phrases for positive sentiment detection. In addition to the lists of positive and negative sentiment phrases based on their POS, there is also an aggregate of words which express positive or negative sentiment in itself, but in context, they often occur with a negation, which results in expressing the opposite sentiment. In the rules, there are 33 negated positive and 17 negated negative words and phrases (an example grammar for detecting negated negative words and phrases is given in Figure 3), which adds up in a total of 203 words and phrases for negative sentiment detection and 146 words and phrases for positive sentiment detection.

sample	precision	recall	F ₁ -score
initial	0.371	0.283	0.321
development	0.435	0.469	0.451
test	0.413	0.393	0.402

Table 3: Polarity phrase detection accuracy of the rule-based component

	+*	-*	x*	precision	recall	F ₁ -score
+	40	3	17	0.677	0.666	0.671
-	2	25	17	0.555	0.568	0.561
x	17	17	30	0.468	0.468	0.468

Table 4: System accuracy on overall sentiment (+, -, x) detection and confusion matrix for overall sentiment assignment (+*, -* and x* represent assignments by the system)

3 Evaluation

The evaluation was conducted on a manually annotated held-out test set containing 11,500 tokens in 168 articles. The initial prototype of the polarity phrase detection module, that was designed from scratch in NooJ, was first evaluated on the test set in a form of a *dry run* test for purposes of further development. The results are given in Table 3 joint for positive and negative polarity phrases. The results of the dry run were shown to be rather low, with an F₁-score of only 0.321. The rules were thus tuned, as previously mentioned, by observing the development set and another two tests were performed with the improved rules — one on the development set itself and the other on the test set. These results are also given in Table 3 and they show an improvement over the baseline for both the development set and the test set. Being that horoscope texts are highly complex in terms of irregularities of phrases, i.e., showing rare re-occurrences of polarity phrases among texts from varying sources, these scores were considered to be a satisfactory entry point for overall article sentiment detection.

The results of system evaluation with respect to overall article sentiment are given in Table 4. The rows of the confusion matrix represent gold standard annotation while the columns present system annotation. The matrix clearly indicates that the system performance is high for the task of discriminating between positive and negative overall sentiment, while its accuracy steeply decreases upon inclusion of the neutral sentiment article category. This observation is also supported by the inter-annotator agreement and the data in Table 1 and 2. The correlation between the number of polarity phrases and overall sentiment given in Table 2 is clearly manifested in the evaluation results, being that the overall performance of the system is satisfactory even if the rule-based phrase detection module performance might be considered somewhat low in absolute terms, especially with respect to those obtained for, e.g., well-structured financial texts (Agić et al., 2010).

Table 4 also shows that positive words and phrases are more accurately detected than the negative ones — the observed difference in F₁-scores of the positive and negative phrase detection is as high as 0.11 in favor of the positive phrase detection. Considering that there are substantially more negative words and phrases in the rules for detection (203 vs. 146) and that there are also considerably more negated positive phrases than vice versa (33 vs. 17),

sign	web sources							+	-	x
aries	x	x	+	x	+	+	x	3	0	4
taurus	-	+	+	+	x	x	x	3	1	3
gemini	+	-	+	-	x	x	x	2	2	3
cancer	-	+	+	x	-	-	x	2	3	2
leo	x	x	x	-	-	x	-	0	3	4
virgo	-	+	+	+	x	+	-	4	2	1
libra	-	-	+	-	+	+	x	3	3	1
scorpio	x	+	x	-	x	-	-	1	3	3
sagittarius	+	+	x	-	-	-	x	2	3	2
capricorn	x	x	+	+	x	x	x	2	0	5
aquarius	+	-	x	-	+	-	+	3	3	1
pisces	+	+	+	+	x	x	x	4	0	3

Table 5: Horoscope sentiment by web source on 2012-05-18

we can conclude that in this type of texts, unlike positive sentiment which is expressed more clearly and explicitly, negative sentiment is often covert and masked with various modifiers and within very complex expressions where negations occur far from the positive word (e.g., in hr. *danas nećete imati baš dobar dan*, en. *you will not have such a good day today*), so they are very difficult to detect with the rules.

Table 5 is an illustration of the sentiment trend information provided by the system. As mentioned previously, the texts are processed on a daily basis and both the texts and the respective annotations are stored in a database. This enables graphical display of sentiment trend across text sources (websites) and text categories (zodiac signs). The table indicates that the overall horoscope sentiment is consistently inconsistent across the seven different web sources and — perhaps even more interestingly — that the possible consistencies might be observed only within single web sources, not respecting the zodiac signs. In the specific case of sentiment analysis in the narrow domain of daily horoscope texts, this might therefore support the claim that perhaps the most reliable sentiment detection feature is the daily sentiment of the text authors. Sentiment trend is more explicitly encoded in Figure 4, as it presents an illustration of a sentiment time series with respect to zodiac signs (top) and web sources of horoscope texts (bottom). Figures for all categories, i.e., zodiac signs, web sources and different time frames are available via the system web interface (<http://lt.ffzg.hr/sentiscope/>).

Conclusion and perspectives

Detecting text sentiment in a very specific and narrow domain such as daily horoscope texts has shown not to be trivial and easy to achieve, given that such texts are characterized both by specific and often very complex phrases and syntax and a particular, domain-dependent style, which can be specific for each individual author, as well. This considered, obtained F_1 -score of 0.566 for overall system accuracy and 0.402 for phrase detection accuracy, with observed annotator agreement of 75.97% (kappa 0.641, linearly weighted kappa 0.593), are here regarded as satisfactory and useful.

For future work, obtained data — the collected texts, the system and the processing results — can be used for different types of linguistic analysis, e.g., discourse analysis and socio-linguistic

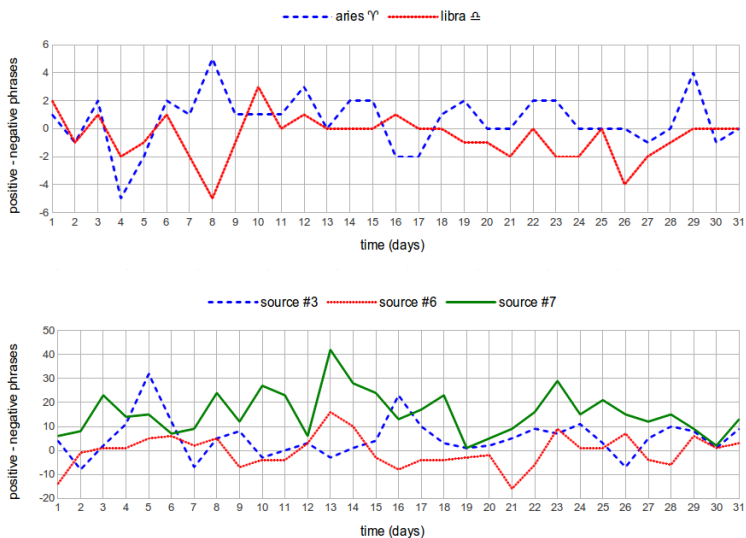


Figure 4: Overall sentiment time series by zodiac sign and web source for March 2012, expressed by the absolute difference between the number of detected positive and negative polarity phrases

analysis. Improvements to the implemented simple link between polarity phrases and overall sentiment might also be investigated, being that the current implementation trivially addresses (especially) neutral sentiment articles. Besides, the developed model could be easily adjusted and applied for sentiment annotation and visualization in other domains.

Acknowledgments

The presented results were partially obtained from research within project CESAR (ICT-PS2 grant 271022) funded by the European Commission, and partially from research within projects 130-1300646-0645 and 130-1300646-1776 funded by the Ministry of Science, Education and Sports of the Republic of Croatia.

References

- Agić Ž, Ljubešić N, Tadić M. (2010). Towards Sentiment Analysis of Financial Texts in Croatian. In *Proceedings of LREC 2010*, ELRA, 2010, pp. 1164–1167.
- Ahmad K, Gillam L, Cheng D. (2005). Society Grids. In *Proceedings of the UK e-Science All Hands Meeting*, Swindon, EPSRC, 2005, pp. 923–930.
- Ahmad K, Gillam L, Cheng D. (2006). Sentiments on a Grid: Analysis of Streaming News and Views. In *Proceedings of LREC 2006*, ELRA, 2006.

- Ahmad K, Cheng D, Almas Y. (2006). Multi-lingual Sentiment Analysis of Financial News Streams. In *Proceedings of the First International Conference on Grids in Finance*, International School for Advanced Studies, Trieste, Italy, 2006.
- Almas Y, Ahmad K. (2007). A note on extracting "sentiments" in financial news in English, Arabic and Urdu. *The Second Workshop on Computational Approaches to Arabic Script-based Languages*, Linguistic Society of America, 2007, pp. 1–12.
- Bruce R, Wiebe J. (1999). Recognizing Subjectivity: A Case Study of Manual Tagging. *Natural Language Engineering*, volume 5, 1999, pp. 187–205.
- Daly N, Kearney C, Ahmad K. (2009). Correlating Market Movements With Consumer Confidence and Sentiments: A Longitudinal Study. *Text Mining Services*, Leipziger Beiträge zur Informatik, 2009, pp. 169–180.
- Devitt A, Ahmad K. (2007). Sentiment Polarity Identification in Financial News: A Cohesion-based Approach. In *Proceedings of the Annual Meeting of the Association of Computational Linguistics (ACL 2007)*, Prague, Czech Republic, 2007.
- Devitt A, Ahmad K. (2008). Sentiment Analysis and the Use of Extrinsic Datasets in Evaluation. In *Proceedings of LREC 2008*, ELRA, 2008.
- Kim S-M, Hovy E. (2006). Identifying and Analyzing Judgment Opinions. In *Proceedings of HLT-NAACL*, 2006.
- Pang B, Lee L. (2008). Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*, volume 2, number 1-2, 2008, pp. 1–135.
- Remus R, Heyer G, Ahmad K. (2009). Sentiment in German language news and blogs, and the DAX. *Text Mining Services*, Leipziger Beiträge zur Informatik, 2009, pp. 149–158.
- Riloff E, Wiebe J, Wilson T. (2003). Learning Subjective Nouns using Extraction Pattern Bootstrapping. In *Proceedings of CoNLL*, 2003, pp. 25–32.
- Shanahan J, Qu Y, Wiebe J. (2006). Computing Attitude and Affect in Text: Theory and Applications. *Information Retrieval Series*, number 20, Springer, 2006.
- Silberstein M. (2004). NooJ: an Object-Oriented Approach. In *INTEX pour la Linguistique et le Traitement Automatique des Langues*, Cahiers de la MSH Ledoux, Presses Universitaires de Franche-Comté, pp. 359–369. See URL <http://www.nooj4nlp.net/>.
- Silberstein M. (2005). NooJ's Dictionaries. In *Proceedings of the Second Language and Technology Conference*, Poznan University, 2005.
- Wiebe J. (2000). Learning Subjective Adjectives from Corpora. In *Proceedings of AAAI*, 2000.
- Wiebe J, Wilson T, Bruce R, Bell M, Martin M. (2004). Learning Subjective Language. *Computational Linguistics*, volume 30, number 3, 2004, pp. 277–308.
- Wiebe J, Bruce R, O'Hara T. (1999). Development and Use of a Gold Standard Data Set for Subjectivity Classifications. In *Proceedings of the 37th ACL Conference*, 1999, pp. 246–253.
- Wilson T, Wiebe J, Hoffmann P. (2005). Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. In *Proceedings of HLT/EMNLP*, 2005, pp. 347–354.

