Creating Subjective and Objective Sentence Classifiers from Unannotated Texts

Janyce Wiebe[1] and Ellen Riloff[2]

[1] Department of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260
`wiebe@cs.pitt.edu`

[2] School of Computing
University of Utah
Salt Lake City, UT 84112
`riloff@cs.utah.edu`

Abstract. This paper presents the results of developing subjectivity classifiers using only unannotated texts for training. The performance rivals that of previous supervised learning approaches. In addition, we advance the state of the art in objective sentence classification, by learning extraction patterns associated with objectivity and creating objective classifiers that achieve substantially higher recall than previous work with comparable precision.

1 Introduction

There has been a recent swell of interest in the automatic identification and extraction of attitudes, opinions, and sentiments in text. Motivation for this task comes from the desire to provide tools for information analysts in government, commercial, and political domains, who want to automatically track attitudes and feelings in the news and on-line forums. How do people feel about recent events in the Middle East? Is the rhetoric from a particular opposition group intensifying? What is the range of opinions being expressed in the world press about the best course of action in Iraq? A system that could automatically identify opinions and emotions from text would be an enormous help to someone trying to answer these kinds of questions. Applications that could benefit from this technology include multi-perspective question answering, which aims to present multiple answers to the user based on opinions derived from different sources, and multi-document summarization, which aims to summarize differing opinions and perspectives.

There is also a need to explicitly recognize objective, factual information for applications such as information extraction and question answering. Linguistic processing alone cannot determine the truth or falsity of assertions, but we could direct the system's attention to statements that are objectively presented, to lessen distractions from opinionated, speculative, and evaluative language.

The goal of our research is to develop learning methods to create classifiers that can distinguish subjective from objective sentences. We strive to develop systems that excel at subjective classification as well as objective classification.

In this paper, we present the results of developing subjectivity classifiers using only unannotated texts for training. The performance of the classifiers rivals that of previous supervised learning approaches to the same task. In addition, we advance the state of the art in objective sentence classification, by learning new objective clues and creating objective classifiers that achieve substantially higher recall than previous work with comparable precision. Our approach begins with a seeding process that utilizes known subjective vocabulary to automatically create training data. This data is then used to train an extraction pattern learner and a probabilistic classifier. Finally, we add a self-training mechanism that improves the coverage of the classifiers, while still relying only on unannotated data.

2   The Data and Classification Task

The texts used in our experiments are English language versions of articles from the world press. The data is from a variety of countries and publications and covers many different topics (it was obtained from the Foreign Broadcast Information Service (FBIS), a U.S. goverment agency). 535 texts from this collection have been manually annotated with respect to subjectivity as part of a U.S. government funded program on automatic question answering.[3] These manually annotated texts comprise the *Multi-Perspective Question Answering (MPQA) corpus* and are freely available at nrrc.mitre.org/NRRC/publications.htm.

The test set used in our evaluations consists of 9,289 of the sentences in the MPQA corpus. None of this test data was used to produce any of the features included in our experiments. 5104 of the sentences in the test set (54.9% of the data) are subjective according to the definitions given below. Thus, the accuracy of a baseline classifier that chooses the most frequent class is 54.9%. Our unannotated text corpus consists of 298,809 sentences from the world press collection, and is distinct from the annotated MPQA corpus.

The annotation scheme and inter-coder reliability studies associated with the MPQA data are described in [1]. The scheme was inspired by work in linguistics and literary theory on *subjectivity*, which focuses on how opinions, emotions, etc., are expressed linguistically in context [2]. The goal is to identify and characterize expressions of *private states* in a sentence. *Private state* is a general covering term for opinions, evaluations, emotions, and speculations [3]. For example, in sentence (1), the writer is expressing a negative evaluation.

(1) "They are no more than frail excuses and pretexts to evade the peace process since this process does not agree with the ideology of expansion and the building of settlements."

---

[3] The ARDA (Advanced Research and Development Activity in Information) AQUAINT (Advanced QUestion and Answering for INTelligence) program.

Sentence (2) reflects the private state of Western countries. Mugabe's use of "overwhelmingly" also reflects a personal private state, his positive reaction to and characterization of his victory.

(2) *"Western countries were left frustrated and impotent after Robert Mugabe formally declared that he had overwhelmingly won Zimbabwe's presidential election."*

The annotators were asked to identify all expressions of private states in each sentence and to indicate various attributes, including strength (*low,medium,high,* or *extreme*). The gold-standard classes used in our evaluations are defined as follows: if a sentence has at least one private state of strength *medium* or higher, then the sentence is subjective; otherwise, it is objective. These are the same definitions that other researchers have used when performing experiments with the MPQA data [4, 5].

3   Learning Subjective and Objective Sentence Classifiers

3.1   Automatically Generating Training Data using Rule-based Classifiers

As a starting point for our research, we reimplemented the high precision, low recall subjective and objective classifiers that we previously developed [4]. We will refer to these as the *rule-based classifiers* because they do not involve learning but merely classify sentences by looking for well-established general subjectivity clues that have been previously published in the literature.[4] Some are drawn from manually developed resources, including entries from [6, 7], Framenet lemmas with frame element *experiencer* [8], and adjectives manually annotated for polarity [9]. Some were learned from corpora, including words distributionally similar to subjective seed words [10], n-grams [11, 12], and subjective nouns learned using extraction pattern (EP) bootstrapping [5]. The clues were divided into strong and weak subjective clues, where strong subjective clues have subjective meanings with high probability, and weak subjective clues have subjective meanings with lower probability.

The rule-based subjective classifier classifies a sentence as subjective if it contains two or more strong subjective clues (otherwise, it does not label the sentence). In contrast, the rule-based objective classifier looks for the absence of clues: it classifies a sentence as objective if there are no strong subjective clues in the current sentence, there is at most one strong subjective clue in the previous and next sentence combined, and at most 2 weak subjective clues in the current, previous, and next sentence combined (otherwise, it does not label the sentence).[5]

---

[4] We will be happy to make these clues available to other researchers.

[5] This is slightly more liberal than in [4], which did not allow a strong subjective clue in the previous or next sentence. This difference explains the higher recall figures reported here.

Our research uses these rule-based classifiers to generate training data for subsequent learning algorithms, which we will describe in the coming sections. Figure 1 shows the first stage of the training data creation process. The rule-based subjective classifier is applied to the unlabeled corpus to identify sentences that it can label as subjective. Similarly, the rule-based objective classifier identifies sentences that it can label as objective. These subjective and objective sentences form our *initial training set*.
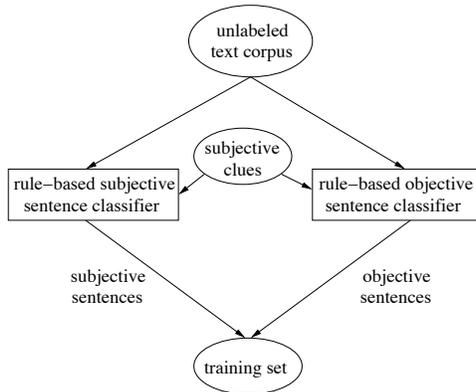

Fig. 1. Initial Training Data Creation

We use the following evaluation metrics in this paper. *Subjective precision (SubjPrec)* is the percentage of sentences automatically classified as subjective that are truly subjective. The *subjective recall (SubjRec)* is the percentage of true subjective sentences that are automatically classified as subjective. The *subjective F-measure (SubjF)* is the usual F-measure combining precision and recall. *ObjPrec*, *ObjRec*, and *ObjF* are defined similarly.

On the annotated test set, the rule-based subjective classifier achieved 34.2% subjective recall and 90.4% subjective precision. The rule-based objective classifier achieved 30.7% objective recall and 82.4% objective precision. Based on these results, we expect that the initial training set generated by these classifiers is of relatively high quality. Of the 298,809 sentences in the unannotated text corpus, the rule-based classifiers labeled 52,918 sentences as subjective and 47,528 as objective, creating a training set of over 100,000 sentences.

3.2 Extraction Pattern (EP) Learning

Previous research has shown that patterns designed for information extraction can effectively represent expressions associated with subjectivity [4]. Objectivity is a different beast because any objective statement can be made subjective by adding a subjective modifier to it. Consequently, it is not clear that individual

expressions can be considered to be truly objective in an absolute sense. However, we hypothesized that in practice there are many expressions that are highly correlated with objective statements and therefore would be strong clues that a sentence is objective. In the Wall Street Journal, for example, sentences containing the words "profits" or "price" are very likely to be objective, even though there is no reason why a subjective sentence could not contain those words.

Consequently, we decided to explore the idea of learning extraction patterns that are correlated with objectivity and then using them as features in a machine learning algorithm. To learn extraction patterns, we used the AutoSlog-TS [13] algorithm because it does not need annotated texts for training. Instead, AutoSlog-TS requires one set of "relevant" texts and one set of "irrelevant" texts. Extraction patterns are created by applying a set of syntactic templates to the corpus. The syntactic constructions recognized by AutoSlog-TS are described in [13] and reflect syntactic relationships identified by a shallow parser.

We trained the EP learner on the initial training set to generate patterns associated with objectivity as well as patterns associated with subjectivity. In our experiments, the subjective sentences were the relevant texts, and the objective sentences were the irrelevant texts. The patterns chosen as the subjective patterns are those that are strongly correlated with subjective sentences, while the patterns chosen as the objective patterns are those that are negatively correlated with subjective sentences (and hence positively correlated with objective sentences). AutoSlog-TS merely ranks patterns in order of their association with the relevant texts, so we automatically selected the best patterns for each class using two thresholds: θ_F is the frequency of the pattern in the corpus, and θ_P is the conditional probability (estimated from the training set) that a text is relevant if it contains the pattern: $\text{Pr}(\text{relevant} \mid pattern_i)$. For our experiments, *subjective patterns* were identified by setting $\theta_F \geq 5$ and $\theta_P \geq .95$ (i.e., at least 95% of its occurrences must have been in subjective sentences). *Objective patterns* were identified by setting $\theta_F \geq 5$ and $\theta_P \leq .15$ (i.e., at most 15% of its occurrences could have been in subjective sentences). Table 1 shows a few examples of subjective and objective patterns that were learned.

Subjective Patterns	Objective Patterns
<subj> believes	<subj> increased production
<subj> was convinced	<subj> took effect
aggression against <np>	delegation from <np>
to express <dobj>	occurred on <np>
support for <np>	plans to produce <dobj>

Table 1. Extraction Pattern Examples

Consider a simple classifier that classifies a sentence as subjective if it contains any of the learned subjective patterns. The subjective precision of this classifier on the manually annotated test set is 74.5% (i.e., 74.5% of the sen-

	SubjRec	SubjPrec	SubjF	ObjRec	ObjPrec	ObjF	Acc
Subj RBC	34.2	90.4	46.6				61.9
Subj RBC w/Patterns	58.6	80.9	68.0				69.7
Obj RBC				30.7	82.4	44.7	65.8
Obj RBC w/Patterns				33.5	82.1	47.6	66.7

Table 2. Rule-Based Classifier Results

tences with subjective patterns are subjective). The subjective recall is 59.8% (i.e., 59.8% of the subjective sentences contain at least one subjective pattern). The similar figures for a cooresponding objective classifier are 71.3% objective precision and 11.7% objective recall (i.e., 71.3% of the sentences with objective patterns are objective, and 11.7% of the objective sentences contain at least one objective pattern). The low objective recall reflects the fact that many fewer instances of objective patterns were found in the data (832 versus 6364 instances of subjective patterns). These results suggest that the EPs are good clues for distinguishing subjective sentences from objective sentences, but are not sufficient by themselves.

Next, we incorporated the learned EPs into the rule-based classifiers as follows. The subjective patterns were added to the set of strong subjective clues, which are used by both the subjective and objective rule-based classifiers. The strategy used by the rule-based subjective classifier remained the same. However, the strategy used by the rule-based objective classifier was augmented as follows: in addition to its previous rules, a sentence is also labeled as objective if it contains no strong subjective clues but at least one objective EP. Note that adding the subjective EPs to the set of strong subjective clues works to *decrease* the recall of the objective classifier because it looks for the absence of subjectivity clues. To balance that effect, the additional test for objective EPs can serve to *increase* the recall of the objective classifier.

The first row of Table 2 shows the results of the original rule-based subjective classifier (*Subj RBC*), and the second row shows the results after adding the subjective extraction pattern clues. Similarly, the third row shows the results for the original rule-based objective classifier (*Obj RBC*), and the fourth row shows the results after adding the objective EP clues. Comparing rows one and two, the subjective precision dropped from 90.4% to 80.9%, but subjective recall increased from 34.2% to 58.6%. Comparing rows three and four, the objective precision decreased only slightly (from 82.4% to 82.1%), and the objective recall increased from 30.7% to 33.5%. Adding EPs to the rule-based classifiers clearly expanded their coverage with relatively smaller drops in precision.

3.3   Naive Bayes Sentence Classification

The labeled sentences identified by the rule-based classifiers provide us with the opportunity to apply supervised learning algorithms to our sentence classification task. Previous work [14, 5, 15] found that naive Bayes performs well for

	SubjRec	SubjPrec	SubjF	ObjRec	ObjPrec	ObjF	Acc
Naive Bayes	70.6	79.4	74.7	77.6	68.4	72.7	73.8

Table 3. Results of Naive Bayes Trained on Initial Training Data

subjectivity recognition, so we used naive Bayes as our learning algorithm. We trained the naive Bayes classifier using the initial training set and several types of set-valued features. There are features for each of the following sets: the strong subjective clues used by the original rule-based classifiers; the weak subjective clues used by the objective rule-based classifier; the subjective patterns generated by the EP learner; and the objective patterns generated by the EP learner. We also added features for the following parts of speech, which were shown to be effective in previous work [5, 15, 11]: pronouns, modals (excluding 'will'), adjectives, cardinal numbers, and adverbs (excluding 'not'). A three-valued feature was defined for each set based on the presence of 0, 1, or ≥ 2 members of that set in the sentence. In addition, to incorporate contextual information in the classifier, another three-valued feature was defined for each set based on the presence of 0, 1, or ≥ 2 members of that set in the previous and next sentences combined.

Row one of Table 3 shows the performance of the naive Bayes classifier on the test set. The classifier achieves relatively balanced recall and precision for both subjective and objective sentences.

3.4 Self-Training the Sentence Classifier

The initial training data used by the naive Bayes classifier was generated by the rule-based classifiers, which simply look for the presence or absence of a set of general subjectivity clues. There are obvious concerns associated with this type of automatically created training data, such as potential biases introduced by the rules. A related concern is that the training sentences will be similar to one another and less heterogenous than the set of sentences that the classifier will ultimately be applied to.

We therefore saw an opportunity to try to improve the classifier by generating a new training set using the classifier itself. The naive Bayes classifier uses a greater variety of features than the rule-based classifiers and it exploits a probabilistic model to make classification decisions based on combinations of these features. We hypothesized that the naive Bayes classifier might be able to reliably label a different, and perhaps more diverse, set of sentences in the unlabeled corpus than the rule-based classifiers did.

The procedure we use is a variant of *self-training*, as the term is used by Nigam and Ghani [16]. They describe the procedure as follows: "Initially, self-training builds a single naive Bayes classifier using the labeled training data and all the features. Then it labels the unlabeled training data and converts the most confidently predicted document of each class into a labeled training example. This iterates until . . ." (p. 90). Rather than adding one instance per

class at a time to a cache of labeled data, we use our naive Bayes classifier to label all the sentences in the entire unannotated corpus from scratch, including those in the initial training set. Then, we select the top $N/2$ most confidently labeled sentences in each class to include in the new training data (where N = the size of the initial training set + 10,000 sentences). The chosen sentences form a brand new training set that we then use to retrain the EP learner and then the naive Bayes classifier. The overall process is depicted in Figure 2.
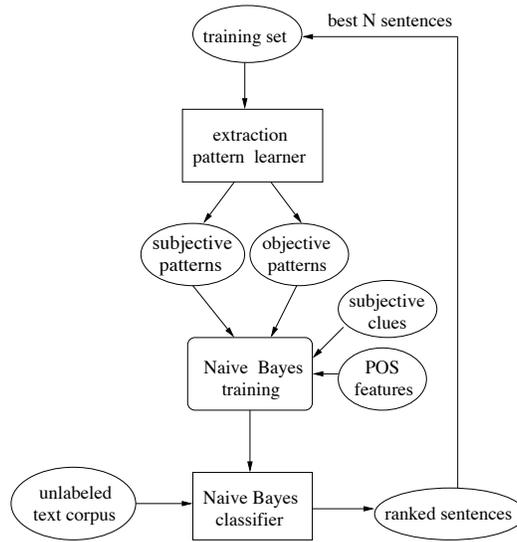

Fig. 2. Self-Training Process

The recall of the learned patterns improved substantially using the new training set, with just a minor drop in precision: subjective precision of the subjective patterns decreased from 74.5% to 73.1%, and objective precision of the objective patterns decreased from 71.3% to 68.9%, while subjective recall of the subjective patterns increased from 59.8% to 66.2% and objective recall of the objective patterns increased from 11.7% to 17.0%.

Table 4 shows the performance of the rule-based classifiers (RBC) and the naive Bayes classifiers on the test set after training on the initial training set (these are the rows labeled 1) and after retraining on the new training set (these are the rows labeled 2).

When the patterns learned on the new training set were incorporated into the rule-based classifiers, the classifiers showed increases in recall but with virtually no drop in precision and even a slight increase for objective sentences (compare rows (a) and (b) for the subjective rule-based classifiers, and rows (c) and (d) for the objective rule-based classifiers).

	SubjRec	SubjPrec	SubjF	ObjRec	ObjPrec	ObjF	Acc
(a) Subj RBC w/Patterns 1	58.6	**80.9**	68.0				69.7
(b) Subj RBC w/Patterns 2	62.4	80.4	70.3				71.0
(c) Obj RBC w/Patterns 1				33.5	82.1	47.6	66.7
(d) Obj RBC w/Patterns 2				34.8	**82.6**	49.0	67.3
(e) Naive Bayes 1	70.6	79.4	74.7	**77.6**	68.4	**72.7**	**73.8**
(f) Naive Bayes 2	**86.3**	71.3	**78.1**	57.6	77.5	66.1	73.4
(g) RWW03 (supervised)	77	81	79	74	70	72	76

Table 4. Comparison of Results

Rows (e) and (f) show that the recall of the naive Bayes classifier swung dramatically toward subjective sentences (+15.7% recall for subjective sentences, -20% recall for objective sentences). At the same time, subjective precision decreased by 8 percentage points while objective precision increased by 9.

Finally, row (g) shows the performance of the best supervised subjectivity sentence classifier on the same type of data [5], which we will denote as RWW03. RWW03 was trained on a subset of the MPQA corpus containing 2197 sentences. 1296 (59%) of those sentences were subjective, so the accuracy of a baseline classifier that chooses the most frequent class was a bit higher for that dataset than for the one used in this paper (its baseline accuracy is 54.9%, as explained in Section 2).

The boldface numbers represent the best results achieved by our classifiers for each evaluation metric. For subjective sentences, the self-trained naive Bayes classifiers achieved the best recall, which was substantially higher than the recall obtained by RWW03, although our precision at the high recall level is lower. The best precision that we obtained is basically the same as RWW03, but with slightly lower recall. For objective sentences, our initial naive Bayes classifier (e) achieved a slightly higher f-measure than RWW03. All in all, our classifiers achieved performance levels comparable to those obtained by a supervised learning system. Our highest precision objective classifier was the rule-based classifier with EPs after self-training (d).

4 Related Work

There has been a recent flurry of research in the related areas of opinion extraction, sentiment analysis, semantic orientation and polarity classification, and subjectivity analysis. Much of this work focuses on lexical aquisition, identifying subjective, positive, or negative words and phrases [9, 17, 9, 18, 10, 19, 20]. Riloff and Wiebe [4] used extraction pattern learning to find subjective expressions, but we know of no previous research on learning objective expressions.

Several projects have focused on document-level subjectivity classification. Some work identifies inflammatory texts (e.g., [21]) or classifies texts as positive or negative ([22, 17, 14]). Research in genre classification has included recognition

of subjective genres such as *editorials* and objective genres such as *business* or *news* (e.g., [23, 24, 12, 15]).

In contrast, our work involves classifying individual sentences. Sentence-level subjectivity classification is useful because most documents contain a mix of subjective and objective sentences. For example, newspaper articles are typically thought to be relatively objective, but [12] reported that, in their corpus, 44% of sentences (in articles that are not editorials or reviews) were subjective.

Almost all previous evaluations of sentence-level subjectivity classifiers involved supervised learning systems (e.g., [15, 5, 11]). We compared our results to [5] in Section 3.4. The precision achieved by [15] was lower, especially for objective sentences. The accuracies reported by [11] are higher (they do not report precision), but their baseline accuracy is very high. However, [15, 11] used different data sets with different annotation schemes, so our results cannot be directly compared.

As described in Section 3.1, [4] report high subjective and objective precisions, but achieve at most 40% subjective recall and 30% objective recall.

Automatic subjectivity/opinion/sentiment analysis is being applied to many interesting applications, including classification of reviews [19, 14, 11, 25, 26], analysis of product reputations [26, 25, 27], tracking sentiments toward events [28, 22, 29], and incorporating opinions into into question answering and multi-document summarization systems [15].

5   Conclusions

We presented the results of developing subjectivity classifiers using only unannotated texts for training. The performance rivals that of previous supervised learning approaches. In addition, we advance the state of the art in objective sentence classification, by learning EPs associated with objectivity and creating objective classifiers that achieve substantially higher recall than previous work with comparable precision.

6   Acknowledgements

References

1. Wilson, T., Wiebe, J.: Annotating Opinions in the World Press. In: Proceedings of the 4th ACL SIGdial Workshop on Discourse and Dialogue (SIGdial-03). (2003) 13–22

2. Banfield, A.: Unspeakable Sentences. Routledge and Kegan Paul, Boston (1982)
3. Quirk, R., Greenbaum, S., Leech, G., Svartvik, J.: A Comprehensive Grammar of the English Language. Longman, New York (1985)
4. Riloff, E., Wiebe, J.: Learning Extraction Patterns for Subjective Expressions. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2003). (2003) 105–112
5. Riloff, E., Wiebe, J., Wilson, T.: Learning Subjective Nouns Using Extraction Pattern Bootstrapping. In: Proceedings of the 7th Conference on Natural Language Learning (CoNLL-2003). (2003) 25–32
6. Levin, B.: English Verb Classes and Alternations: A Preliminary Investigation. University of Chicago Press, Chicago (1993)
7. Ballmer, T., Brennenstuhl, W.: Speech Act Classification: A Study in the Lexical Analysis of English Speech Activity Verbs. Springer-Verlag (1981)
8. Baker, C., Fillmore, C., Lowe, J.: The Berkeley FrameNet Project. In: Proceedings of the COLING-ACL. (1998)
9. Hatzivassiloglou, V., McKeown, K.: Predicting the Semantic Orientation of Adjectives. In: Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL-97). (1997) 174–181
10. Wiebe, J.: Learning Subjective Adjectives from Corpora. In: Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000). (2000) 735–740
11. Dave, K., Lawrence, S., Pennock, D.M.: Mining the Peanut Gallery: Opinion Extraction and Semantic Classification of Produce Reviews. In: Proceedings of the 12th International World Wide Web Conference (WWW2003). (2003) Web Proceedings.
12. Wiebe, J., Wilson, T., Bell, M.: Identifying Collocations for Recognizing Opinions. In: Proceedings of the ACL-01 Workshop on Collocation: Computational Extraction, Analysis, and Exploitation. (2001) 24–31
13. Riloff, E.: Automatically Generating Extraction Patterns from Untagged Text. In: Proceedings of the Thirteenth National Conference on Artificial Intelligence, The AAAI Press/MIT Press (1996) 1044–1049
14. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? Sentiment Classification Using Machine Learning Techniques. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2002). (2002) 79–86
15. Yu, H., Hatzivassiloglou, V.: Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinion Sentences. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-2003). (2003) 129–136
16. Nigam, K., Ghani, R.: Analyzing the Effectiveness and Applicability of Co-Training. In: Proceedings of the Ninth International Conference on Information and Knowledge Management. (2000) 86–93
17. Turney, P.: Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-2000). (2002) 417–424
18. Hatzivassiloglou, V., Wiebe, J.: Effects of Adjective Orientation and Gradability on Sentence Subjectivity. In: 18th International Conference on Computational Linguistics (COLING-2000). (2000)
19. Turney, P., Littman, M.: Measuring Praise and Criticism: Inference of Semantic Orientation from Association. ACM Transactions on Information Systems (TOIS) **21** (2003) 315–346

20. Gordon, A., Kazemzadeh, A., Nair, A., Petrova, M.: Recognizing Expressions of Commonsense Psychology in English Text. In: Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics (ACL-03). (2003) 208–215

21. Spertus, E.: Smokey: Automatic Recognition of Hostile Messages. In: Proceedings of the Eighth Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-97). (1997) 1058–1065

22. Das, S.R., Chen, M.Y.: Yahoo! for Amazon: Opinion Extraction from Small Talk on the Web. In: Proceedings of the 8th Asia Pacific Finance Association Annual Conference. (2001)

23. Karlgren, J., Cutting, D.: Recognizing Text Genres with Simple Metrics Using Discriminant Analysis. In: Proceedings of the Fifteenth International Conference on Computational Linguistics (COLING-94). (1994) 1071–1075

24. Kessler, B., Nunberg, G., Schütze, H.: Automatic Detection of Text Genre. In: Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL-97). (1997) 32–38

25. Nasukawa, T., Yi, J.: Sentiment Analysis: Capturing Favorability Using Natural Language Processing. In: Proceedings of the 2nd International Conference on Knowledge Capture (K-CAP 2003). (2003)

26. Morinaga, S., Yamanishi, K., Tateishi, K., Fukushima, T.: Mining Product Reputations on the Web. In: Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2002). (2002)

27. Yi, J., Nasukawa, T., Bunescu, R., Niblack, W.: Sentiment Analyzer: Extracting Sentiments about a Given Topic using Natural Language Processing Techniques. In: Proceedings of the 3rd IEEE International Conference on Data Mining (ICDM-2003). (2003)

28. Hearst, M.: Automatic Acquisition of Hyponyms from Large Text Corpora. In: Proc. of the 14th International Conference on Computational Linguistics (COLING-92). (1992)

29. Tong, R.: An Operational System for Detecting and Tracking Opinions in Online Discussions. In: Working Notes of the SIGIR Workshop on Operational Text Classification. (2001) 1–6