Creating a Fuzzy Believer to Model Human Newspaper Readers

Ralf Krestel[1], René Witte[1], and Sabine Bergler[2]

[1] Institut für Programmstrukturen und Datenorganisation (IPD)
Universität Karlsruhe (TH), Germany
[2] Department of Computer Science and Software Engineering
Concordia University, Montréal, Canada

Abstract. We present a system capable of modeling human newspaper readers. It is based on the extraction of reported speech, which is subsequently converted into a fuzzy theory-based representation of single statements. A domain analysis then assigns statements to topics. A number of fuzzy set operators, including fuzzy belief revision, are applied to model different belief strategies. At the end, our system holds certain beliefs while rejecting others.

1 Introduction

With the huge success of the internet, the natural language processing (NLP) research community has developed whole branches that deal explicitly with vast amounts of unstructured information encoded in written natural language. One goal is to gain knowledge about hard facts like "The number of inhabitants of city X" or the "name of the president of country Y." But a lot of information, especially within newspaper articles, are not hard facts, which could be easily proven right or wrong. Often newspaper articles contain different *views* of the same event, or state controversial opinions about a certain topic. In this case the notion of *belief* becomes relevant.

For humans, this is a daily task. Depending on context information and background knowledge, together with other belief structures, humans tend to believe certain statements while other statements are rejected. The process of believing also varies between different humans, not only depending on their different background knowledge, but also on different attitudes towards a coherent worldview or importance and their ability of logic thinking.

For a computational system simulating a human newspaper reader by imitating his belief processing, this involves not only the extraction of beliefs stated in an article, but also their comparison to existing beliefs held by the system. Such an *artificial believer* [1] must have different belief strategies to model different human approaches.

The application area of an artificial believer is large. Potential users include:

– Companies interested in customers' opinions about their own products or products from a competitor.
– Governments interested in the opinions of people about their country or the goverment's work.

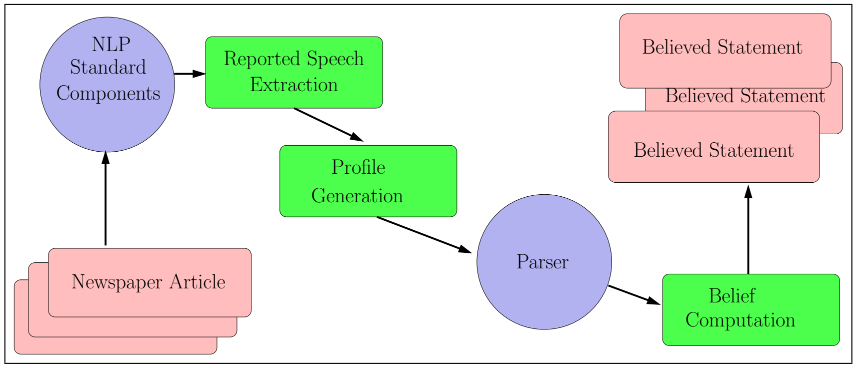

Fig. 1. Fuzzy Believer System Architecture Overview

– Individuals, who wish to have a personalized news digest compiled automatically.

Our system is designed for the last group of users, but is not limited to this application.

The system we present in this paper addresses various problems within the NLP domain. Our main contributions are: 1. Developing rules to identify and extract reported speech from newspaper articles; 2. processing the gained information by applying fuzzy set theory to natural language processing; 3. creating a working implementation of these ideas, together with an evaluation environment.

The remainder of this paper is structured as follows: In the next section, we give an overview of our fuzzy believer system, followed by a more detailed description of the individual components in Section 2. An evaluation of our approach, using different corpora and evaluation methods, is presented in Section 3. Section 4 discusses related work, followed by conclusions in Section 5.

2 Design and Implementation

The core concept embodied in our approach is the application of fuzzy set theory to the NLP domain. This allows for an explicit modeling of fuzziness inherent to natural languages and enables the user to control the system's behaviour by varying various runtime parameters responsible for the fuzzy processing. Reported speech statements present the basic set of beliefs for our system. These kinds of statements usually express a belief held by the source of the statement and allows a clear attribution of the statement to this source. The extracted reported speech structures are further processed and the output of external semantic parsers is utilized to identify predicate-argument structures (PAS) within the reported speech content. Each PAS defines a statement, which the system eventually either believes or rejects. They also form the foundation for the fuzzy processing and the basis for our heuristics to process beliefs.

To mirror the different processing steps, our fuzzy believer system consists of a set of components running consecutively. It is implemented using GATE (General Architecture for Text Engineering) [2], which offers a framework for developing NLP applications. For preprocessing, we use a number of standard components shipped with GATE, for high-level processing we developed our own components. An overview of the system's structure is shown in Figure 1, indicating the four main components constituting our system: (1) Reported speech extraction; (2) Profile generation; (3) PAS extraction; (4) Belief computation.

2.1 Previous Work

A similar system extracting reported speech from newspaper articles together with its source and reporting verb is presented in [3] and [4]. The system passes the extracted information through evidential analysis and processes the results to different *profiles*.

In detail, to evolve profiles out of *basic profiles*, which consist of a statement and its source, an intermediate step *(merged profiles)* is needed. In this step, the exploitation of coreference information becomes necessary. For this reasen, a noun phrase coreferencer [5] is used to identify same sources of different statements. These statements are then merged into a single merged profile.

Evidential chains are generated and a percolation algorithm is used, see [6]. The merged statements are grouped according to the reporter who uttered the reported speech. This allows to model different degrees of confidence into a certain newspaper, a certain reporter, and a certain source. To encode the different confidences in the resulting profile, a dichotomy of held beliefs and potential beliefs is introduced.

In contrast to our fuzzy believer, this system is limited to handling beliefs without considering their content, solely based on information about the source of a reported clause and the reporter of the article. Apart from improving the *extraction of reported speech*, the system presented in this paper is capable of *identifying the topic* of the reported speech and for each topic the *polarity of individual statements* concerning the topic. On top of this information, an *artificial believer* is implemented simulating knowledge acquisition through different strategies.

2.2 Extracting Reported Speech

The main source for our fuzzy believer stems from *reported speech* in newspaper articles. This allows us to explicitly attribute statements to sources. Additional information that can be analyzed and therefore has to be identified by the extracting component comprise the reporting verb and circumstantial information.

To find reported speech structures, we identified six patterns around 50 verbs [7] that are often used within reported speech constructs. Example 1 shows a typical reported speech structure and identifies the different elements.

$$\underbrace{\overbrace{\text{Last October}}^{\text{circumstances}},\ \overbrace{\text{his brother Hubert}}^{\text{source}}\ \overbrace{\text{told}}^{\text{rverb}}\ \overbrace{\text{the bankruptcy court}}^{\text{addressee}}}_{\text{reporting clause}}\ \text{that}\ \underbrace{\text{Paul was very ill}}_{\text{reported clause}}.\text{[3]} \tag{1}$$

This information can be utilized to perform evidential analysis [8], thereby assigning different degrees of *confidence* in a statement according to the reliability of the source and the reporting verb used.

[3] sentence from Wall Street Journal 03.03.88

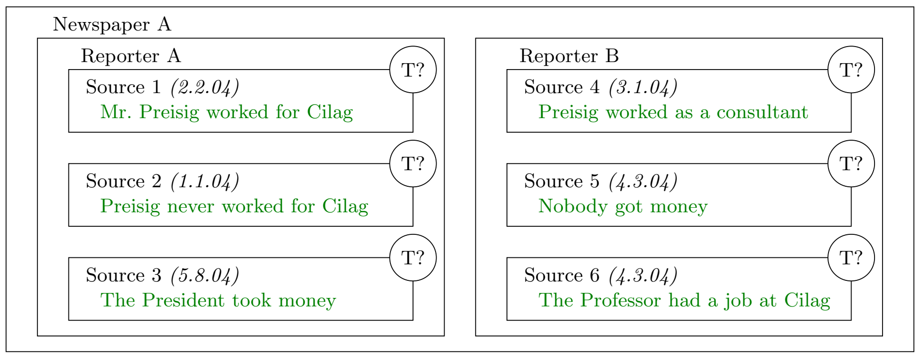

Fig. 2. Information after extracting reported speech – sources are isolated and topics (T) not yet identified

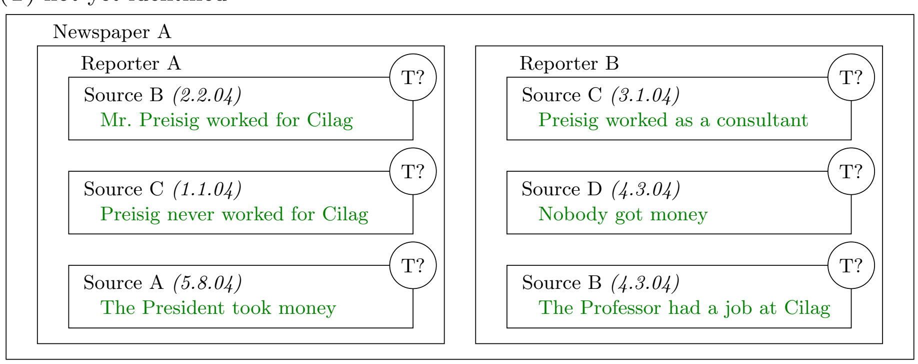

Fig. 3. The different statements after identifying the source entities

Figure 2 shows the results of the reported speech extraction component assuming 6 fictitious newpaper articles[4] dealing with two different topics. We adapted a presentation scheme for beliefs proposed by Ballim and Wilks [1], using nested boxes to visualize the held beliefs of different actors. Each box contains a statement together with its source and the publishing date. Every statement is assigned to the reporter who wrote the article containing the statement, and finally the newspaper who published the article is named.

2.3 Generating Profiles

The profile generation component assembles the reported speech fragments and prepares them for the next processing step. A profile assigns each statement to a source, reporter, and newspaper. Basically, the component extracts the reported speech clauses, which can then be further processed by a parser. It also adds coreference information for each source by traversing the data structure created by our fuzzy coreference resolution system [5]. Figure 3 shows our example with the added coreference information.

[4] inspired by articles in WSJ 12.03.86

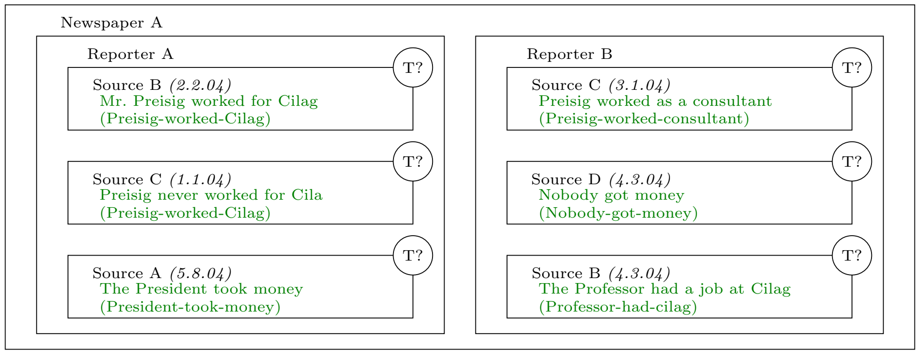

Fig. 4. The different extracted predicate-argument structures

2.4 Extracting Predicate-Argument Structures

To decide whether a sentence has the same topic as another one, we need to find a way to compare sentences with each other. To facilitate this task, we do not compare whole sentences, but their predicate-argument structures, consisting of "subject," "verb," and "object." Because one sentence might contain more than one statement, a correct syntactic analysis is paramount for predicate-argument structure (PAS) generation. Our experiments showed that no single parser is consistently reliable enough for PAS extraction. Thus, our PAS extraction component can work with the results of three different parsers: RASP [9], MiniPar [10], and SUPPLE [11].

The PAS extractor applies a custom rule set for each of these parsers in order to determine subject, verb, and object of a statement.

The extracted predicate-argument structures for our example can be seen in Fig. 4. To demonstrate the system, we chose rather simple sentences containing only one PAS each but the algorithm can handle more complex structures as well.

2.5 Computing Beliefs

The core of our system is the *fuzzy believer* component. Its tasks are:

1. Identify a topic for each statement.
2. Compute the fuzzy representation for each statement to identify polarity.
3. Process fuzzy information for each topic according to a strategy.
4. Generate a graphical view of the result.

Identifying Domains. The first step is to group the statements into *domains* according to their topics. These domains constitute the basic sets for the fuzzy operations performed later on; basically, they partition the statement space into individual domains, which can be processed independently. Every domain represents one topic identified by the extracted PASs.

To determine if a statement fits into an existing domain, we use *heuristics* to measure the semantic proximity of each new statement with the statements

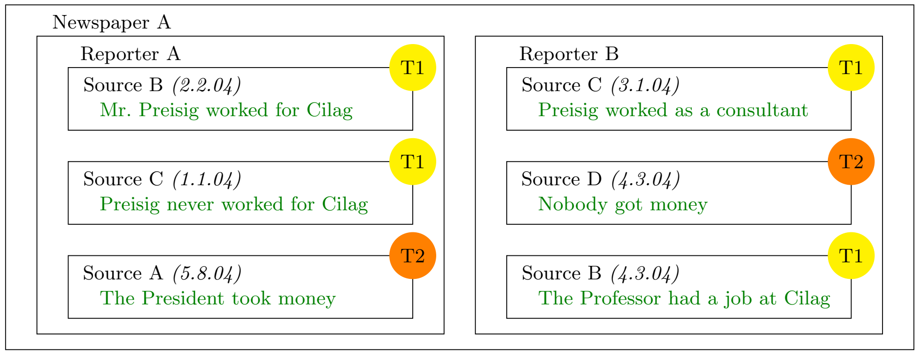

Fig. 5. The different topics after identifying the domains

in all existing domains. For this, the system applies two main heuristics: (1) A WordNet [12] related heuristic, and (2) a substring heuristic.

These heuristics compare the PAS elements of one statement with the elements of the other statements in one domain and return a value representing how similar the heuristics consider the two PAS elements. A runtime option defines if *strict* matching is neccessary to include a new statement in a domain, or if a more *lenient* matching is sufficient. For a strict match, the new statement's PAS must be similar to all existing statements within a domain. In case of a lenient match, the new statement needs only to be similar to one statement of a domain, essentially implementing a transitive relation on the domain elements.

To cause a match between two statements, at least two parts of their corresponding PAS structures must be similar enough. That means, the value assigned by a heuristic must exceed the defined threshold for either subject and object, subject and verb, or verb and object.

This approach permits assigning a statement to more than one domain. If a new statement does not fit into any of the existing domains, a new domain is dynamically created, initially containing this statement.

Each domain contains all statements that have the same or opposite meaning. In other words, we try to identify each fact in the world and arrange all statements concerning this fact in one domain. The example in Fig. 5 should contain two domains after the classification process: *T1* "Someone taking money" and *T2* "Preisig working as consultant at Cilag." The different statements are assigned a label identifying the topic.

Identifying Polarity. In the next step, the statements gathered for each domain have to be evaluated by identifying their polarity. The goal is to identify opposing statements by using different fuzzy heuristics. The fuzzy representation μ_{S_i} of a statement S_i contains the degrees of similarity of this statement with all other statements within the same domain. Each degree is normalized to a fuzzy value in the $[0, 1]$-interval and can be interpreted as the semantic distance between two statements. Fig. 6 shows the fuzzy representation of a statement S_1 within a domain containing five statements ($S_1, \ldots, S_5$). The fuzzy sets are interpreted in a possibilistic fashion: A fuzzy value of 0 indicates no possible semantic similarity

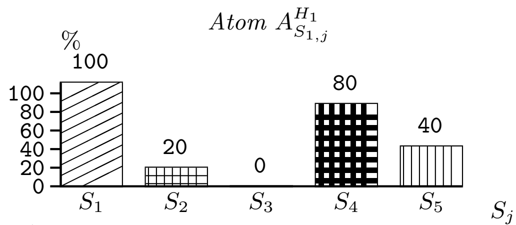

Fig. 6. Statement $A_{S_{1,j}}$ with correlation grades for all statements in the domain $(S_1, \ldots, S_5)$ as computed by heuristic H_1

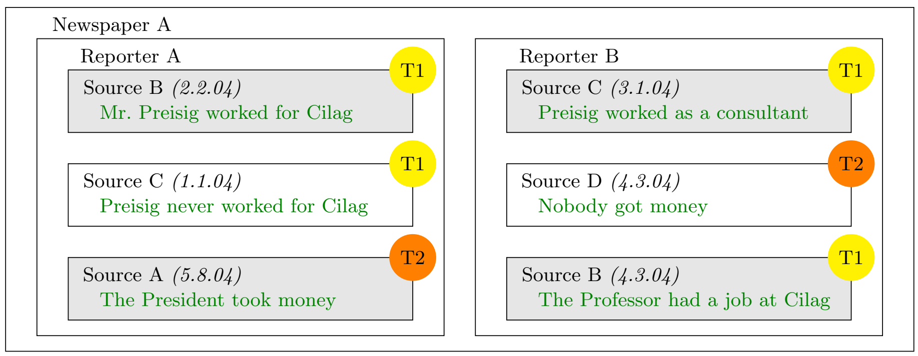

Fig. 7. Believe Majority: The system believes statements with grey background.

between the two statements, while a value of 1.0 indicates the highest possibility of similarity between them. In the current implementation, only one heuristic is used. It compares the verbs of two statements using their WordNet semantic distance to find synonyms and antonyms.

Computing Beliefs. One of the crucial parts of the fuzzy believer system is to decide which of the collected statements to believe and which to reject. For this, each domain is processed independently. The system implements the following strategies: (1) Believe majority, (2) believe old news, (3) believe new news, (4) believe certain source/reporter/newspaper, and (5) believe weighted majority. The strategies are based on fuzzy processing. Three fuzzy operations are essential to implement the strategies: *Merging*, *expanding*, and *revising*. These operations are computed directly on the fuzzy set representation of each statement, which has been generated as described above.

Based on the fuzzy representation, the merge operation groups all statements into one class, if a threshold of semantic similarity is reached. Usually, merging all statements leads to two classes within each domain, one containing statements about a topic and the other one containing opposing statements about this topic. The majority strategy picks the class with the most statements and marks them as belief. Fig. 7 shows the result of this strategy for our example. For topic T2 there is no majority, in this case the system chooses either of the statements.

The expansion operator initially believes the first statement in a domain and each new statement becomes included only if it is compatible with all the ones existing in a domain. Expansion [13] can be used to implement the "Believe old news" strategy by ordering the processing according to the publishing date.

For the fuzzy belief revision process [14], new statements are always believed and only those of the existing statements that are not in conflict with the new ones are kept. This is exactly what we need for the "Believe new news" strategy. For the weighted majority strategy, we use information of the majority strategy and combine it with information about the reliability of the newspaper, the reporter, and the source of the statement. To believe in a certain source, reporter, or newspaper, fuzzy processing is not necessary and this strategy can be implemented utilizing the profile generator information.

2.6 Summary

Our fuzzy believer system processes natural language articles and identifies the topics discussed in a text. Statements are extracted from the texts based on reported speech structures and assigned to domains, which form the formal basis for automatic processing using fuzzy operators. The main believer component can simulate different reading strategies, like a reader accepting all new information (and erasing conflicting old knowledge), or a stubborn reader clinging to old beliefs while rejecting all incompatible new information.

The output of the fuzzy believer system is a set of held beliefs and rejected beliefs acquired from "reading" a document collection. Presently, we export this result into a graphical representation using the LaTeX-format similar to the presentation of the examples in this paper.

3 Evaluation

The system we present here is complex and attempts a novel analysis. Therefore no Gold standard corpora are available. In order to evaluate our system we have thus chosen to evaluate its components separately on standard reference resources in related domains.

Extracting Reported Speech. In order to evaluate the reported speech extraction component, we randomly picked 7 newspaper articles from the Wall Street Journal corpus. The articles contain about 400 sentences (∼6100 words), among them 133 reported speech constructs. For the detection of reporting verb and source, our system achieved a recall value of 0.83 and a precision value of 0.98. This results in an f-measure of 0.90.

Identifying Domains. The domain finding task is quite hard and error-prone. Remember that the domain classification is solely based on the predicate-argument structures extracted from the output of one of the three deployed parsers. The evaluation of the domain finding component includes a comparison of the results obtained with RASP, MiniPar, and manually annotated predicate-argument structures (gold standard). The conservative strategy of SUPPLE, which only marks relations that are considered to be 100% correct, proved to be not applicable, as it creates too few extractable PAS.

The test data we used is taken from the MSR corpus [15] and comprised 300 paraphrase pairs. We assumed all sentences were reported clauses to skip the

Table 1. Domain Classification: Recall and Precision for different parsers

Configuration	*Recall*			*Precision*		
	Rasp	Minipar	Manual	Rasp	Minipar	Manual
3-3-3-lenient	0.59	0.54	0.56	0.57	0.63	0.78
3-3-3-strict	0.59	0.50	0.55	0.63	0.75	0.85
5-5-5-lenient	0.70	0.60	0.62	0.29	0.39	0.29
5-5-5-strict	0.52	0.52	0.52	0.41	0.53	0.54
5-3-5-lenient	0.65	0.51	0.59	0.31	0.57	0.45
5-3-5-strict	0.59	0.58	0.52	0.56	0.41	0.61

reported speech extraction part from distorting the domain finding evaluation results. The special layout of the test corpus, containing pairs of paraphrases and thus two statements per topic, made it necessary to develop a method to measure the performance accurately. The fact that one sentence can contain more than one statement, represented as different predicate-argument structures, made the evaluation scenario more complex. We conducted a test with a test set of 116 paraphrase pairs, which was additionally annotated by hand with predicate-argument structures. This allows an estimation on the influence of the parser and the parser extraction component on the domain classification process. The results can be found in Table 1.[5]

Identifying Polarity. To test the sense detection or opinion grouping function, we would need a special corpus containing test data with opposing and supporting statements for a special opinion, which are semantically close enough to fulfill the requirement of belonging to the same domain. The data that comes closest to these conditions are the entailment pairs of the PASCAL challenge corpus [16]. There are some minor drawbacks, though.

Firstly, the positive entailment examples are rather easy to evaluate, because if one sentence entails another, the senses of the two sentences must have the same direction. But non-entailment between two sentences doesn't necessary imply opposing opinions in these sentences. But fortunately this is often the case for the PASCAL-2 challenge corpus we used. A second problem is the fact that sentence pairs, especially without entailment, would not be assigned the same domain by our domain classification algorithm, therefore it is not possible to evaluate the data using only the polarity identification component.

We solved these problems by checking the non-entailing examples manually for opposing sentences and developing a scheme to measure the performance of the sense detection algorithm without influence from the domain finding component. This scheme comprises the consideration of only those statement pairs that are correctly assigned to the same domain.

We tested different configurations and computed accuracy for two different settings. For one experiment, we included all results in the evaluation, counting

[5] The configuration settings in the table mean, from left to right: Maximum WordNet Distance between (1) subjects, (2) verbs, (3) objects of two statements. And (4) indicates whether a new statement has to match with one (lenient) or all (strict) statements within one domain.

Table 2. Polarity Identification: Accuracy values for different parse methods

Configuration	*Accuracy*			
	Sense & Domain		*Only Sense*	
	Rasp	Minipar	Rasp	Minipar
3-3-3-strict-0.7	0.52	0.55	0.53	0.58
5-5-5-lenient-0.7	0.51	0.53	0.51	0.53
5-5-5-strict-0.3	0.52	0.53	0.55	0.51
5-5-5-strict-0.7	0.51	0.54	0.50	0.56
7-7-7-strict-0.7	0.51	0.52	0.51	0.52

the entailment pairs that were not grouped into the same domain by the domain classification as non-entailing. In the table, this is refered to as "Sense & Domain." The other test setting only considered the sentence pairs that were actually grouped into the same domain by the domain classification component. That way, we limited the influence of the domain classification algorithm on the sense detection. An overview of the achieved perfomance is shown in Table 2[5] with the additional configuration parameter showing the threshold for assigning the same polarity to a statement.

4 Related Work and Discussion

The extraction of opinions from newspaper articles [17] or customers reviews [18, 19] has become an active research field. Those approaches are usually only concerned with the identification and extraction of information without processing it further, except for binary classification within a clearly specified domain.

In the wake of the PASCAL challenge [16,20], systems have been developed to deal with the relation of sentences to each other. The different approaches include the recognition of false entailement [21], or learning entailement [22]. Others are concerned with relatedness between words and how to measure it [23]. We were not interested in concentrating on one of these areas but rather to develop an all-embracing system incorporating different aspects.

The results our system achieved for extracting reported speech is highly competitive. Doandes [24], using a different subset of the WSJ-corpus, reports a recall of 0.44 and a precision of 0.92 for their system compared to 0.83 and 0.98 our system obtained.

For the domain classification, our best results for 300 paraphrase pairs from the MSR-corpus are: Precision 38%, Recall *81%* and Precision *52%*, Recall 58%. These values can probably be improved by using more sophisticated heuristics, although there will be a ceiling set by the parser and by the use of language in general. The same meaning can be expressed by various different sentences whose words are not in close relations to each other and therefore hard to detect by current NLP tools. Keeping these facts in mind, the obtained numbers are rather satisfactory and promising for future development.

The rather shallow semantic approach sets a practical limit to the achievable results. This can be infered by comparing the numbers obtained using manually parsed predicate-argument structures with the numbers obtained by the parsers. It shows that there is space for improvement on the side of the parsers, as well as on the side of the PAS extractor. Combining the results of different parsers could also lead to better results, but a precision of 55% and a recall of 85%, as obtained for the best configuration of the system using manually parsed PASes, shows that it needs more and/or better heuristics to get a really significant improvement.

The polarity identification task was expectedly the hardest one. This is illustrated by the rather poor results we obtained by trying to find different opinions within one domain. Best accuracy values were obtained using Minipar and were around 58%. This task is very hard for computational systems. But with more elaborated heuristics it is possible to increase these numbers, comparable to the Pascal challenge [16, 20], where systems also started with around 50% accuracy and improved over time.

Testing of the different strategies revealed that the fuzzy processing operators perform in accordance to their assigned tasks. Further evaluation of the results would need some kind of measure to get quantitative, comparable results. This is beyond the scope of this paper and deferred to future work.

5 Summary and Conclusions

We presented a fuzzy believer system, which is capable of differentiating between different topics and different polarity of statements and decides what to believe based on configurable strategies. The system was applied to processing reported speech information, generating a belief set containing the knowledge obtained from "reading" different newspaper articles. Our approach is based on the application of fuzzy set theory to natural language processing resulting in a fuzzy believer with variable belief strategies.

The results for the individual subtasks are promising but the development of a measure to evaluate the system as a whole is still pending. The growing number of available news sources, blogs, and webpages makes it necessary to facilitate the information gathering for humans. Our fuzzy believer was designed to deal with huge amounts of information and supports a user's opinion finding process.

References

1. Ballim, A., Wilks, Y.: Artificial Believers: The Ascription of Belief. Lawrence Erlbaum Associates, Inc. (1991)
2. Cunningham, H., Maynard, D., Bontcheva, K., Tablan, V.: GATE: A framework and graphical development environment for robust NLP tools and applications. In: Proc. of the 40th Anniversary Meeting of the ACL. (2002)
3. Bergler, S.: Conveying attitude with reported speech. In Shanahan, J.C., Qu, Y., Wiebe, J., eds.: Computing Attitude and Affect in Text: Theory and Applications. Springer Verlag (2005)
4. Bergler, S., Doandes, M., Gerard, C., Witte, R.: Attributions. [25] 16–19
5. Witte, R., Bergler, S.: Fuzzy Coreference Resolution for Summarization. In: Proc. of 2003 Intl. Symposium on Reference Resolution and Its Applications to Question Answering and Summarization (ARQAS), Venice, Italy (June 23–24 2003) 43–50
6. Gerard, C.: Modelling Readers Of News Articles Using Nested Beliefs. Master's thesis, Concordia University, Montréal, Québec, Canada (2000)

7. Quirk, R.: A comprehensive grammar of the English language. Longman Group Limited (1985)
8. Bergler, S.: The Evidential Analysis of Reported Speech. PhD thesis, Brandeis University, Massachusetts, USA (1992)
9. Briscoe, E., Carroll, J., Watson, R.: The Second Release of the RASP System. In: Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions. (2006)
10. Lin, D.: Dependency Based Evaluation of MINIPAR. In: Proc. of the Workshop on the Evaluation of Parsing Systems, First Intl. Conf. on Language Resources and Evaluation. (1998)
11. Gaizauskas, R., Hepple, M., Saggion, H., Greenwood, M.A., Humphreys, K.: SUPPLE: A practical parser for natural language engineering applications. In: Proceedings of the 9th International Workshop on Parsing Technologies (IWPT2005), Vancouver (2005)
12. Fellbaum, C., ed.: WordNet: An Electronic Lexical Database. MIT Press (1998)
13. Witte, R.: Architektur von Fuzzy-Informationssystemen. BoD (2002) ISBN 3-8311-4149-5.
14. Witte, R.: Fuzzy Belief Revision. In: 9th Intl. Workshop on Non-Monotonic Reasoning (NMR'02), Toulouse, France (April 19–21 2002) 311–320
15. Dolan, B., Brockett, C., Quirk, C.: Mirosoft research paraphrase corpus. Online: www (march 2005) `http://research.microsoft.com/research/nlp/msr_paraphrase.htm`.
16. Bar-Haim, R., Dagan, I., Dolan, B., Ferro, L., Giampiccolo, D., Magnini, B., Szpektor, I.: The Second PASCAL Recognising Textual Entailment Challenge. In: Proc. of the Second PASCAL Challenges Workshop on Recognising Textual Entailment. (2006)
17. Bethard, S., Yu, H., Thornton, A., Hatzivassiloglou, V., Jurafsky, D.: Automatic extraction of opinion propositions and their holders. [25] 20–27
18. Gamon, M., Aue, A., Corston-Oliver, S., Ringger, E.K.: Pulse: Mining customer opinions from free text. In Famili, A.F., Kok, J.N., Peña, J.M., Siebes, A., Feelders, A.J., eds.: Advances in Intelligent Data Analysis VI, 6th International Symposium on Intelligent Data Analysis, IDA 2005, Madrid, Spain, September 8-10, 2005, Proceedings. Volume 3646 of Lecture Notes in Computer Science., Springer (2005) 121–132
19. Kim, S.M., Hovy, E.: Identifying and analyzing judgment opinions. [26] 200–207
20. Dagan, I., Glickman, O., Magnini, B.: The PASCAL Recognising Textual Entailment Challenge. In: Proc. of the PASCAL Challenges Workshop on Recognising Textual Entailment. (2005)
21. Snow, R., Vanderwende, L., Menezes, A.: Effectively using syntax for recognizing false entailment. [26] 33–40
22. MacCartney, B., Grenager, T., de Marneffe, M.C., Cer, D., Manning, C.D.: Learning to recognize features of valid textual entailments. [26] 41–48
23. Klebanov, B.B.: Measuring Semantic Relatedness Using People and WordNet. [26] 13–16
24. Doandes, M.: Profiling For Belief Acquisition From Reported Speech. Master's thesis, Concordia University, Montréal, Québec, Canada (2003)
25. Qu, Y., Shanahan, J., Wiebe, J., eds.: Exploring Attitude and Affect in Text: Theories and Applications. Technical Report SS-04-07. AAAI Press, Stanford, CA, USA (March 22–25 2004)
26. Moore, R.C., Bilmes, J., Chu-Carroll, J., Sanderson, M., eds.: Proceedings of the Human Language Technology Conference of the NAACL, Main Conference. Association for Computational Linguistics, New York City, USA (June 2006)