

# Negation scope detection with a conditional random field model<sup>①</sup>

Lydia Lazib, Zhao Yanyan, Qin Bing<sup>②</sup>, Liu Ting

(Research Center for Social Computing and Information Retrieval, Harbin Institute of Technology, Harbin 150001, P. R. China)

## Abstract

Identifying negation cues and their scope in a text is an important subtask of information extraction that can benefit other natural language processing tasks, including but not limited to medical data mining, relation extraction, question answering and sentiment analysis. The tasks of negation cue and negation scope detection can be treated as sequence labelling problems. In this paper, a system is presented having two components: negation cue detection and negation scope detection. In the first phase, a conditional random field (CRF) model is trained to detect the negation cues using a lexicon of negation words and some lexical and contextual features. Then, another CRF model is trained to detect the scope of each negation cue identified in the first phase, using basic lexical and contextual features. These two models are trained and tested using the dataset distributed within the \*Sem Shared Task 2012 on resolving the scope and focus of negation. Experimental results show that the system outperformed all the systems submitted to this shared task.

**Key words:** negation detection, negation cue detection, negation scope detection, natural language processing

## 0 Introduction

Negation, as simple as it can be in concept, is a complex and an essential phenomenon in any language. It has the ability to inverse the meaning of an affirmative statement into its opposite meaning. In a sentence, the presence of negation is indicated by the presence of a negation cue. The negation cue is a lexical element that carries negation meaning. A negation cue can occur in different forms<sup>[1]</sup>: as an explicit negation, which can be a single word negation (e.g., “no”, “not”) or a multiple words negation (e.g., “neither...nor” or “rather than”), as an implicit negation, where syntactic patterns imply negative semantics (e.g., “This movie was below my expectations.”), or as a morphological negation, where word roots are modified with a negating prefix (e.g., “un-”, “in-” or “dis-”) or negating suffix (e.g., “-less”). The scope of negation is the sequence of words in the sentence that is affected by the negation cue<sup>[2]</sup>. For example, in Sentence (1) the word *not* is the negation cue, and the discontinuous word sequences ‘*Holmes*’ and ‘*say anything*’ form the scope.

[Holmes] did not [say anything]. (1)

For many NLP (natural language processing) ap-

plications, distinguishing between affirmative and negative information is an important task. A system that does not deal with negation would treat the facts in these cases incorrectly as positives. For example, in sentiment analysis detecting the negation is a critical process, as it may change the polarity of a text and results in a wrong prediction. And in query answering systems failing to account for negation can result in giving wrong answers.

However, most of the systems developed for processing natural language data do not consider the negation present in the sentence. Although, various works have dealt with the identification of negations and their scope in sentences, machine learning techniques started to be used since the creation of the Bioscope corpus<sup>[3]</sup>. This corpus boosted several research experiments on scope resolution. Ref. [4] proposed a supervised system that found the negation cue and their scopes in biomedical texts. The system consists of two memory-based engines, one decides if the tokens in a sentence are negation cues, and another finds the full scope of these negation cues. Ref. [2] used a decision tree to classify whether a token is at the beginning, inside or outside a negation cue. In the scope finding task, they use three classifiers (k-nearest neighbor,

① Supported by the National High Technology Research and Development Programme of China (No. 2015AA015407), the National Natural Science Foundation of China (No. 61273321), and the Specialized Research Fund for the Doctoral Program of Higher Education (No. 20122302110039).

② To whom correspondence should be addressed. E-mail: bing.qin@gmail.com  
Received on May 12, 2016

SVM and CRF classifier) to predict whether a token is the first token in the scope sequence, the last or neither. And Ref. [5] proposed a new approach for tree kernel-based scopes detection by using structured syntactic parse information. Their experiments on the Bioscope corpus showed that both constituent and dependency structured syntactic parse features have the advantage in capturing the potential relationship between cues and their scopes. But the processing of these feature-rich methods takes a lot of effort and knowledge, and is time-consuming.

Negation has also been studied in the context of sentiment analysis, as a wrong interpretation of a negated fact may lead to assign a wrong polarity to the text, and results in a wrong prediction. Ref. [6] affirmed that applying a simple rule that considers only the two words following the negation keyword as being negated by that keyword to be effective. This method yielded a significant increase in overall sentiment classification accuracy. Ref. [1] proposed a negation detection system based on CRF, modeled using features from an English dependency parser. And report the impact of accurate negation detection on state-of-the-art sentiment analysis system. Their system improved the precision of positive sentiment polarity detection by 35.9% and negative sentiment polarity detection by 46.8%. Ref. [7] proposed a sophisticated approach to identify negation scopes for Twitter sentiment analysis. They incorporated to their sentiment classifier several features that benefit from negation scope detection. The results confirm that taking negation into account improves sentiment classification performance significantly. And Ref. [8] used machine learning methods to recognize automatically negative and speculative information, and incorporated their approach to a sentiment classifier. The results achieved demonstrate that accurate detection of negations is of vital importance to the sentiment classification task.

In this paper, a system is proposed to detect the negation cues and their corresponding scopes, in respect to the closed task of the \* Sem Shared Task 2012<sup>[9]</sup>. The system is divided into two cascade sub-tasks, one that detects the negation cues, and another that detects the scope of negation cues identified by the first sub-task. A CRF model is trained using lexical and contextual features on both sub-tasks. These features, compared to those used in previous work<sup>[5,10-12]</sup> are simple to process and are less time-consuming.

This paper includes four sections. In Section 1 the approach is presented to solve the two tasks of the system. Section 2 describes the corpus used in the experiment. Then, in Section 3 different experimental set-

tings, the experimental results and the corresponding analysis are presented. And finally, concluding remarks are contained in Section 4.

## 1 System description

The system is decomposed to identify the scope of negation into two cascade sub-tasks: negation cue detection and negation scope detection. The scope detection is dependent on the task of finding the negation cues.

To conduct the experiment, the corpus provided by organizers of the \* Sem Shared Task 2012 (will be described in Section 2) is used. In order to do a fair comparison with the work submitted during this shared task, all the requirements of the task are followed, which means that the classifier can use only the information provided in the training dataset, without using any external tool to respect the requirement of the closed task.

The approach to detect the negation cues and their respective scope in a sentence is to consider the two tasks as sequence labelling problems. The conditional random fields (CRF)<sup>[13]</sup> classifier has proved in several previous work to be more efficient in solving this kind of problems<sup>[10-12]</sup> compared with other machine learning techniques.

The negation cues present in the test data are identified by training a CRF model using some lexical and contextual features and a lexicon of negation cues that are present in the training data. To identify the scope of negation, a CRF model is also trained using lexical and contextual features.

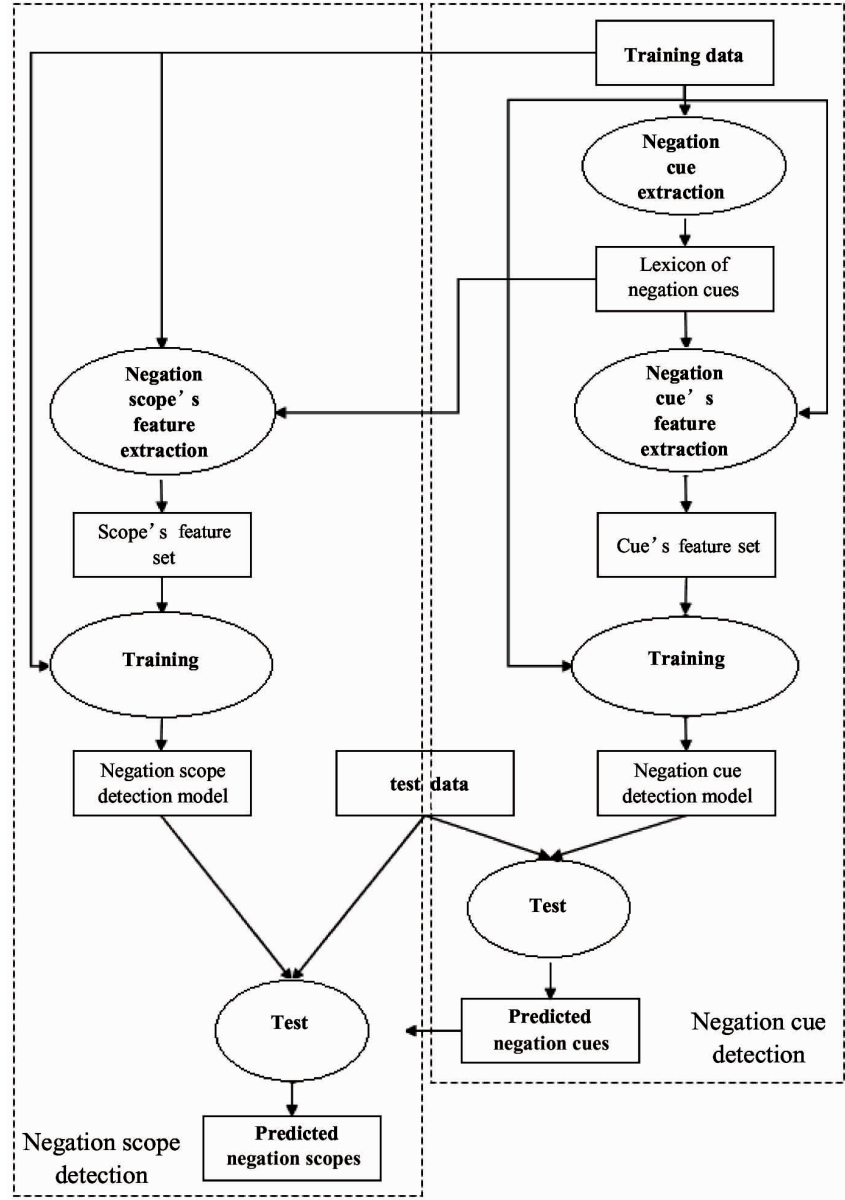
The illustration of the approach is shown in Fig. 1.

The details of the two sub-tasks are described in the subsections below.

### 1.1 Negation cue detection

In this task, all the negation cues presented in the sentences are identified. The negation cues are grouped under five types, depending upon how they are present in the data. They can be:

- Single word cues: not, no, never, etc.
- Continuous multiword cues: rather than, by no means, etc.
- Discontinuous multiword cues: neither... nor, etc.
- Prefix cues: words starting with the prefixes: in-, im-, un-, etc. e. g. impossible.
- Suffix cues: words ending with the suffix: -less, e. g. useless.



**Fig. 1** Illustration of the whole approach of the proposed negation cue and scope detection

Because of the restriction of the closed task of the shared task, any external dictionary or lexicon for negation words couldn't be used. The system creates two lexicons, one corresponding to the negation words (single word, continuous multiword, and discontinuous multiword) by collecting all the negation words that can appear in the training dataset, and another corresponding to the different negation affixes (prefixes, suffixes) found in the training dataset. These two lists are shown in Table 1 and Table 2.

The CRF model was trained on a set of lexical and contextual features. The process of features selection started by combining different features used in some previous work similar to the task<sup>[10,11]</sup>. This feature set has been polished from the features that was proved

empirically to be useless, and added new features that improved the cue detection task. The final feature set used to train the CRF model is presented in Table 3.

Table 1 Lexicon of negation cues (single and multi-words)		
no	not	* n' t
never	nothing	none
except	prevent	nowhere
without	refuse	absence
lack	fail	rather than
nobody	neither ... nor	nothing at all
on the contrary	by no means	not for the world

Table 2 Lexicon of negation cues ( prefixes and suffixes )

in-	im-
dis-	un-
il-	non-
ir	-less

Table 3 Feature set for negation cue detection

Feature name	Description
Token	The word as it appears in the sentence
Lemma	The lemmatized form of the token
POS	The Part-of-speech tag of the token
is _ Punct	The token is a punctuation mark
StartsWithPrefix	The token starts with a prefix negation
EndsWithSuffix	The token ends with a suffix negation
is _ SubToken	The token is a sub-token of a prefix negation, e. g. “possible” is the sub-token of “impossible”
Match _ Single	The token is a single negation word
Multi _ Cont	The token is part of a continuous multiple word negation
Multi _ Disc	The token is part of a discontinuous multiple word negation

The CRF model used considers the features of the current token, two previous and two forward tokens. It also uses features conjunctions by combining features of neighboring tokens, and bigram features.

The CRF model can classify a word in a sentence as being:

- “O”: not part of any negation cue
- “NEG”: as a single word negation
- “NEG \_ DIS”: as a discontinuous multiword negation
- “NEG \_ CONT”: as a continuous multiword negation
- “PRE”: as a prefix negation
- “SUF”: as a suffix negation

The system uses the lexicon presented in Table 1 to set the values of the three features: “Match \_ Single”, “Multi \_ Continuous” or “Multi \_ Discontinuous”, by searching all the words in the sentences of the data that match the words in the lexicon. For each word that matches the system sets the corresponding feature to 1 and set the others to 0.

The system uses the lexicon presented in Table 2 the same way, except that here it searches whether the words in the sentence start with a prefix, end with a suffix or are just simple words. The features modified in this case are: “StartsWithPrefix” and “EndsWith-Suffix”. If the current token is detected as a prefix (or suffix) negation cue, a special treatment should be

done on this token to split it into a prefix (or suffix) negation cue and its corresponding sub-token, using a simple regular expression method.

1.2 Negation scope detection

A CRF model is also trained to identify the scope of negation. The negation cues identified in the first sub-task are used as the new cues, and the scope of which will be identified. If a sentence contains more than one negation cue, then each one will be treated separately, by creating a new training/test instance of the same sentence for each negation cue presenting in the same sentence.

To identify the scope, the model considers the features shown in Table 4. These features are token specific features (Token) and contextual features (e.g., relative position of the token to the cue, whether the token and the cue are in the same segment, etc.), and have the advantage in capturing the potential relationship between cues and their scopes (e.g. the number of token between the current token and the cue, the relative position of the current token from the cue: before, after or same, etc.), and make the prediction of the scope more relevant. They are also simple features, which are less time-consuming, comparing to other features used in other methods<sup>[10-12]</sup>. New features are also created by combining neighboring features and bigram features.

Table 4 Feature set for negation scope detection

Feature name	Description
Token	The word as it appears in the sentence
is _ Cue	If the current token is a negation cue
Distance	The number of token between the current token and the cue
Same _ seg	If the current token and the cue are in the same segment (no punctuation between them)
Relative _ pos	The relative position of the current token from the cue: before, after or same
Cue _ subToken	If the current token starts with a prefix negation (im-, in-, un-, etc.), or ends with a suffix negation (-less)

The CRF model can classify a token in a sentence as being inside (I) or outside (O) the scope of negation. And if the negation cue starts or ends with one of the affixes (prefixes or suffixes) listed in Table 2, the scope of negation includes only the part of the negation cue excluding the affix. Thus, for each negation cue having these affixes, the affix is removed from the cue

and the remaining part is considered as the scope.

2
Data set

Research on negation scope detection has mainly focused on the biomedical domain, and has been neglected in open-domain because of the lack of corpora. One of the only freely available corpus is the dataset released by the organizers of the \* Sem Shared Task 2012. This dataset includes stories of Conan Doyle, and is annotated with negation cues and their corresponding scope, as well as the event that is negated. The cues are the words that express negation, and the scope is the part of the sentence that is affected by the negation cue. The negation event is the main event actually negated by the negation cue. Table 5 shows an example sentence from the dataset.

Table 5
Example sentence annotated for negation following \* Sem Shared Task 2012

File	Sentence#	Token#	Token	Lemma	POS	Parse tree	Cue1	Scope1	Event1	Cue2	Scope2	Event2
WL2	108	0	After	After	IN	(S(S(PP *	-	After	-	-	-	-
WL2	108	1	his	his	PRPMYM	( NP *	-	his	-	-	-	-
WL2	108	2	habit	habit	NN	* ) )	-	habit	-	-	-	-
WL2	108	3	he	he	PRP	( NP *	-	he	-	-	-	-
WL2	108	4	said	Say	VBD	( VP *	-	said	said	-	-	-
WL2	108	5	nothing	nothing	NN	( NP * ) ) )	nothing	-	-	-	-	-
WL2	108	6	,	,	,	*	-	-	-	-	-	-
WL2	108	7	and	and	CC	*	-	-	-	-	-	-
WL2	108	8	after	after	IN	( S( PP *	-	-	-	-	after	-
WL2	108	9	mine	mine	NN	( NP * ) )	-	-	-	-	mine	-
WL2	108	10	I	I	PRP	( NP *	-	-	-	-	I	-
WL2	108	11	asked	ask	VBD	( VP *	-	-	-	-	asked	asked
WL2	108	12	no	no	DT	( NP *	-	-	-	no	-	-
WL2	108	13	questions	question	NNS	* ) ) )	-	-	-	-	questions	-
WL2	108	14	.	.	.	* )	-	-	-	-	-	-

The annotation of cues and scopes is inspired by the Bioscope corpus annotation<sup>[3]</sup>, but there are some differences. The first difference is that the cue is not considered to be part of the scope. The second, the scopes can be discontinuous, and includes the subject, which is not the case in the Bioscope corpus. And finally, the morphological negations are annotated as in the example in Sentence (2) below:

[He] declares that he heard cries but [is] un[able to state from what direction they came]. (2)

Statistics for the corpus are presented in Table 6. More information about the annotations guidelines are provided by Ref. [14], including inter-annotator agreement.

- Column 1: contains the name of the file
- Column 2: contains the sentence number within the file
- Column 3: contains the token number within the sentence
- Column 4: contains the token
- Column 5: contains the lemma
- Column 6: contains the part-of-speech ( POS) tag
- Column 7: contains the parse tree information
- Column 8 to last:
  - If the sentence does not contain a negation, column 8 contains “ \*\*\* ” and there are no more columns.
  - If the sentence does contain negations, the information for each one is provided in three columns: the cue, a word that belongs to the scope, and the negated event, respectively.

Table 6
Corpus statistics

	Training	Development	Test
# tokens	65450	13566	19216
# sentences	3644	787	1089
# negation sentences	848	144	235
% negation sentences	23.27	18.29	21.57
# cues	984	173	264
# scopes	887	168	249

3
Experiments and Results

The CRF model for negation cue and scope detection is trained and tested against the datasets described in Section 2. The identification of cues and scopes is evaluated using the evaluation tool provided by the or-

ganizer of the \*Sem Shared Task 2012 which uses the standard precision, recall, and F-measure metrics to evaluate the system. The evaluation is performed on different levels:

- 1) Cue level: the metrics are performed only for the cue detection.
- 2) Scope CM (cue match): the metrics are calculated at scope level, but require a strict cue match. All tokens of the cue have to be correctly identified.
- 3) Scope tokens (no cue match): the metrics are performed at the token level. The total of scope tokens in a sentence is the sum of tokens of all scopes. For example, if a sentence has 2 scopes, one with 5 tokens and another with 4, the total number of scope tokens is 9.

The punctuation marks are ignored by the evaluation tool to relax the scope evaluation.

The CRF++ tool is used to train the two CRF models for negation cue and negation scope detection.

The results obtained by the system over the test data are shown in Table 7.

	Precision (%)	Recall (%)	F1 (%)
Cues	94.56	92.80	93.69
Scope CM (cue match)	90.96	64.66	75.59
Scope tokens (no cue match)	89.70	83.54	86.51

The analysis of the results obtained by each component are described in the subsections below.

### 3.1 Negation cue detection

The system achieves an F1 score of 93.69% in the task of negation cue detection using a CRF model. The effectiveness of the system is limited by the coverage of the lexicon. Due to the low coverage of the lexicon, the system fails to identify negation cues that are present only in the test data and never appear in the training data. However, the results still outperform all the results submitted to the shared task. As can be seen in Table 9, the system outperforms the system of the participant FBK<sup>[11]</sup>, who got the first place on cue detection task in the shared task.

Also, some words such as *never*, *nothing*, *not*, *no* and *without* are mostly present as negation cues in the data, but not always. Such in the phrase *no doubt*, which is present nine times in the test data, but the word *no* is a negation cue in only four of them. The word *save* is also present once as a negation cue in the training data, but is never a negation cue in the test data. Therefore, our system invariably predicts these occurrences of *save* in the test data as negation cues.

### 3.2 Negation scope detection

The system is able to achieve an F1 score of 86.51% for negation scope detection on scope tokens level (without cue match), and an F1 score of 75.59% on scope level (with cue match). The results show that our system has a higher precision than recall in identifying the scope. As mentioned earlier, the negation cues identified in the first task are used to identify the scope of negation. Using a test data with a 7% error in negation cues as the input to this component and some of the errors of the system in predicting the scope lead to a low recall value in the scope detection.

Table 8 shows the results of the negation scope detection system using the gold cues. These results demonstrate the effectiveness of our system on scope detection task, with an increase of almost 5% on the scope level (with cue match) and 3% on scope tokens level (without cue match).

	Precision (%)	Recall (%)	F1 (%)
Scope CM (cue match)	98.82	67.47	80.19
Scope tokens (no cue match)	91.52	86.75	89.07

The results are compared with the three best work results submitted for the \*Sem shared task 2012 in Table 9. The system outperforms these three best work on both cue and scope detection tasks. The participant FBK<sup>[11]</sup> also used a CRF model to identify the negation cues in a sentence, but has omitted to use the feature related to the “Token”, which is considered as the most valuable feature to identify the negation cue. For scope detection (on scope level), the participant UWashingt<sup>[12]</sup> used essentially lexical and syntactic features, but has neglected the features that capture the relationship between the current token and the negation

		Precision (%)	Recall (%)	F1 (%)
Cue	FBK <sup>[11]</sup>	93.41	91.29	92.34
	Ours	94.56	92.80	93.69
Scope (no cue match)	UWashingt <sup>[12]</sup>	82.90	64.26	72.40
	Ours	90.96	64.66	75.59
Scope tokens (no cue match)	UiO1 r1 <sup>[15]</sup>	81.99	88.81	85.26
	Ours	89.70	83.54	86.51

cue in a sentence. And for scope detection (on scope tokens level), the participant UiO1<sup>[15]</sup> used the SVM classifier to identify the scope of negation. Their results are outperformed by more than 3%, which emphasizes the theory that the CRF classifier is more effective in resolving sequence labelling problems.

## 4 Conclusion

In this paper an approach is proposed to identify the negation cue and the scope of negation in a sentence. It is shown that considering these two tasks as sequence labelling problems, and using CRF model to solve them achieves a considerable accuracy. However, the system cannot cover the negation cues that are not present in the training data. It also misclassifies some negation cues that can appear in non-negated contexts. Moreover, in order to improve the overall accuracy of the scope detection, an accurate system is needed to detect the negation cues, since the errors in the negation cue detection propagates to the identification of the scope. Using features that capture the relationship between the tokens and the negation cue are relevant in identifying the scope of negation.

As future work, we would like to use an extensive lexicon of negation cues to better predict the negation cues. It is also intended to use the current system to solve some problems related to the negation. It is believed that this kind of system can improve the accuracy of several work that are sensitive to the polarity in the information extraction and natural language processing domain, like sentiment analysis or question answering systems.

## References

- [1] Councill I G, McDonald R, Velikovich L. What's great and what's not: learning to classify the scope of negation for improved sentiment analysis. In: Proceedings of the Workshop on Negation and Speculation in Natural Language Processing, ACL, Uppsala, Sweden, 2010. 51-59
- [2] Morante R, Daelemans W. A metalearning approach to processing the scope of negation. In: Proceedings of the 13th Conference on Computational Natural Language Learning, Boulder, USA, 2009. 21-29
- [3] Szarvas G, Vincze V, Farkas R, et al. The BioScope corpus: annotation for negation, uncertainty and their scope in biomedical texts. In: Proceedings of the Workshop on Current Trends in Biomedical Natural Language Processing, Columbus, USA, 2008. 38-45
- [4] Morante R, Liekens A, Daelemans W. Learning the scope of negation in biomedical texts. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, Honolulu, USA, 2008. 715-724
- [5] Zou B, Zhou G, Zhu Q. Tree Kernel-based Negation and Speculation Scope Detection with Structured Syntactic Parse Features, In: Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, Seattle, USA, 2013. 968-976
- [6] Hogenboom A, Van Iterson P, Heerschop B, et al. Determining negation scope and strength in sentiment analysis. In: Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC), Kyoto, Japan, 2011. 2589-2594
- [7] Reitan J, Faret J, Gambäck B, et al. Negation scope detection for twitter sentiment analysis. In: Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA), Lisbon, Portugal, 2015. 99-108
- [8] Cruz N P, Taboada M, Mitkov R. A machine learning approach to negation and speculation detection for sentiment analysis. *Journal of the association for information science and technology*, 2016, 67(9): 2118-2136
- [9] Morante R, Blanco E. \*SEM 2012 shared task: Resolving the scope and focus of negation. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the 6th International Workshop on Semantic Evaluation, Montreal, Canada, 2012. 265-274
- [10] Abu-Jbara A, Radev D. UMichigan: A conditional random field model for resolving the scope of negation. In: Proceedings of the 1st Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the 6th International Workshop on Semantic Evaluation, Montreal, Canada, 2012. 328-334
- [11] Chowdhury M, Mahbub F. FBK: Exploiting phrasal and contextual clues for negation scope detection. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the 6th International Workshop on Semantic Evaluation, Montreal, Canada, 2012. 340-346
- [12] White J P. UWashingon: Negation resolution using machine learning methods. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the 6th International Workshop on Semantic Evaluation, Montreal, Canada, 2012. 335-339
- [13] Lafferty J, McCallum A, Pereira F C N. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of the 18th International Conference on Machine Learning (ICML), Massachusetts, USA, 2001. 282-289
- [14] Morante R, Schrauwen S, Daelemans W. Annotation of negation cues and their scope: Guidelines v1. *Computational linguistics and psycholinguistics technical report series*, CTRS-003, 2011
- [15] Read J, Velldal E, Øvrelid L, et al. UiO1: Constituent-based discriminative ranking for negation resolution. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the 6th International Workshop on Semantic Evaluation, Montreal, Canada, 2012. 310-318

**Lydia Lazib**, born in 1990. She is a Ph.D candidate in Harbin Institute of Technology. She received her B.S. and M.S. degrees in Computer Science Department of Mouloud MAMMERI University of Tizi-Ouzou, Algeria in 2011 and 2013 respectively. Her research interests include sentiment analysis and negation detection.