

The Telling Tail: Signals of Success in Electronic Negotiation Texts

Marina Sokolova

IRO, Université de Montréal
Montréal, Québec, Canada

sokolovm@iro.umontreal.ca

Vivi Nastase

EML Research, gGmbH
Heidelberg, Germany

nastase@eml-research.de

Stan Szpakowicz

SITE, University of Ottawa
Ottawa, Ontario, Canada

ICS, Polish Academy of Sciences
Warsaw, Poland

szpak@site.uottawa.ca

Abstract

We analyze the linguistic behaviour of participants in bilateral electronic negotiations, and discover that particular language characteristics are in contrast with face-to-face negotiations. Language patterns in the later part of electronic negotiation are highly indicative of the successful or unsuccessful outcome of the process, whereas in face-to-face negotiations, the first part of the negotiation is more useful for predicting the outcome. We formulate our problem in terms of text classification on negotiation segments of different sizes. The data are represented by a variety of linguistic features that capture the gist of the discussion: negotiation- or strategy-related words. We show that, as we consider ever smaller final segments of a negotiation transcript, the negotiation-related words become more indicative of the negotiation outcome, and give predictions with higher accuracy than larger segments from the beginning of the process.

1 Introduction

We use language every day to convince, explain, manipulate and thus reach our goals. This aspect of language use is even more obvious in the context of negotiations. The parties must reach an agreement on the partitioning or sharing of a resource, while each party usually wants to leave the negotiation table with the larger piece of the pie. These tendencies become stronger when negotiators use only

electronic means to communicate, that is to say, participate in electronic negotiations. In face-to-face contact, prosody and body language often have a crucial role in conveying attitudes and feelings. E-negotiators, on the other hand, must rely only on texts. We perform automatic analysis of the textual data in e-negotiations. We identify linguistic expressions of such negotiation-specific behaviour that are indicative of the final outcome of the process – success or failure – and observe how powerful a tool language is in helping people get what they want.

In this paper we focus on the negotiation as an ongoing process. We analyze the linguistic features of messages exchanged at various points in the course of the negotiation, to determine the time frame in which the outcome becomes decided. From our experimental point of view, we determine the segment of the negotiation which is most predictive of the outcome. There is an imposed three-week deadline in the electronic negotiations that we analyze. We hypothesize that the pressure of the deadline is reflected in the messages exchanged. The messages written later in the process are more indicative of the outcome of the process. Our empirical results support this hypothesis; an analysis of the linguistic features that make this prediction possible shows what the negotiators' main concerns are as the deadline draws near.

Here is what our results contribute to the field of text analysis. Research on text records of face-to-face negotiations suggests that the language patterns used in the first half of a negotiation predict the negotiation outcome better than those in the second half (Simons, 1993). The explanation was that

in the first phase people establish contact, exchange personal information and engage in general polite conversation, creating a foundation of trust between partners. No numerical data, however, supported this diagnosis, and there was no distinction between the prediction of successful and unsuccessful outcomes. When it comes to text classification, our hypothesis says that the classification of the second parts of e-negotiation texts is more accurate with respect to the outcome than the classification of the first parts. This makes e-negotiation texts different from newsgroup messages, newspaper articles and other documents classified by Blatak et al. (2004), where texts showed better classification accuracy on their initial parts. We report the results of several sets of Machine Learning (ML) experiments. Performed on varying-size text data segments, they support our hypothesis.

We worked with a collection of transcripts of negotiations conducted over the Internet using the Web-based negotiation support system *Inspire* (Kersten and Zhang, 2003). Kersten and Zhang (2003) and Nastase (2006) classified e-negotiation outcomes using non-textual data. Classification based on texts is discussed in (Sokolova et al, 2005; Sokolova and Szpakowicz, 2006). None of those experiments considered segmenting the data, although Sokolova and Szpakowicz (2006) analyzed the importance of the first part of e-negotiations. The work we present here is the first attempt to investigate the effect of parts of e-negotiation textual data on classification quality. In this study we do not report types of expressions that are relevant to success and failure of negotiations. These expressions have been presented and analyzed in (Sokolova and Szpakowicz, 2005).

In section 2 we take a brief look at other work on the connection between behaviour and language. In section 3 we present our data and their representation for ML experiments, and we further motivate our work. Section 4 describes the experiments. We discuss the results in Section 5. Section 6 draws conclusions and discusses a few ideas for future work.

2 Background Review

Young (1991) discusses the theory that the situation in which language is used affects the way in which it

is used. This theory was illustrated with a particular example of academic speech.

The field of neuro-linguistic programming investigates how to program our language (among other things) to achieve a goal. In the 1980s, Rodger Bailey developed the Language and Behaviour Profile based on 60 meta-programs. Charvet (1997) presents a simplified approach with 14 meta-programs. This profile proposes that people's language patterns are indicators of behavioural preferences. In the study of planning dialogues (Chu-Carroll and Carberry, 2000), Searle's theory of speech acts used through the discourse analysis also supports the fact that language carries much of people's behaviour and emotions. Reitter and Moore (2007) studied repetitions in task-oriented conversations. They demonstrated that a speaker's short-term ability to copy the interlocutor's syntax is autonomous from the success of the task, whereas long-term adaptation varies with such success.

We consider a negotiation to be a communication in which the participants want to reach an agreement relative to the splitting/sharing of resources. Language is one of the tools used to reach the goal. We propose that not all messages exchanged throughout a negotiation have the same effect on the negotiation outcome. To test this hypothesis, we take an ever smaller segment of the negotiation, and see how well we can predict the outcome of the process, based only on the messages in this fragment.

We encountered several challenges in predicting e-negotiation outcomes using the messages exchanged. First, electronic negotiations usually do not have a sequential-stage model of behaviour (Koeszegi et al, 2007), which is common in face-to-face negotiations (Adair and Brett, 2005). Here is an example of behavioural phases in face-to-face negotiations: *Perform Relational Positioning* → *Identify the Problem* → *Generate Solutions* → *Reach Agreement*. Unexpected turns and moves – typical of human behaviour – make prediction of the negotiation outcome difficult. In case of electronic negotiation, the absence of the usual negotiation structure further complicates the outcome prediction. This differs e-negotiations from agent-customer phone conversations studied in (Takeuchi et al, 2007), where an agent follows call flow pre-defined by his company's policy.

The longer an e-negotiation takes, the more elaborate the structure of the e-negotiation process becomes. Simpler e-negotiation may involve an exchange of well-structured business documents such as pre-defined contract or retail transactions. A more complex process comprises numerous offers and counter-offers and has a high degree of uncertainty because of the possible unpredictability of negotiation moves.

The next challenge stems from the limitations imposed by the use of electronic means. This overloads text messages with various tasks: negotiation issues themselves, introductions and closures traditional in negotiations, and even socializing. On the other hand, electronic means make the contacts less formal, allowing people to communicate more freely. As a result, the data have a high volume of informality such as abbreviations or slang.

One challenge is specific to text analysis. E-negotiations usually involve a multi-cultural audience of varied background, many of whom are not native English speakers. While communicating in English, they introduce a fair amount of spelling and grammatical mistakes.

We employ various ML methods to determine whether the last segment of the negotiation is better at predicting the outcome, and how small it can be still to perform well.

3 Textual Data in Electronic Negotiations

Participants in a negotiation assume well-defined roles (such as buyers/sellers in some business negotiations, or facilitators in legal disputes), have goals, and adopt specific behaviour to achieve those goals (Koeszegi et al, 2007). These circumstances are reflected in the language of texts exchanged in negotiations, and distinguish this type of texts from casual e-mail exchange and postings on discussion groups and chat boards. We claim that the language captured in e-negotiation textual data changes as a negotiation progresses, and that this is clearly detectable, even though it does not follow a sequential-stage model common in face-to-face-negotiations (Adair and Brett, 2005). To support this hypothesis, we have conducted a series of ML experiments on negotiation segments of varying size and position, using the largest available data of electronic negotiations.

Our data come from the Web-based negotiation support system *Inspire*. *Inspire* has been used in business courses to teach students about e-negotiations and give them a chance to practice bilateral business negotiations conducted in a lightly controlled environment. For many users, conducting negotiations has been a course assignment. Some users wanted to develop their English skills by participating in an *Inspire*-enabled negotiation. A negotiation would last up to three weeks, after which, if an agreement has not been reached, the systems would terminate the negotiation and record it as unsuccessful. The following is an example of a negotiation message (we kept the original spelling):

Dear Georg, I hope you are doing well. I send you this message to ask you what happened to our offer. Just be aware that we will not be indefinitely waiting on your response. As I told you during our last meeting, Itex Manufacturing needs a partnership. So it is important to me to know if you are ready to negotiate with us. We can not afford losing so much precious time. We give you now five more days to answer our offer (1st of december 1997, 2400 Swiss time). After this dead line, will propose our services to your concurrence. I still believe in a good partnership and relationship between our two societies. Let me know if you think so. For Itex Manufacturing. Rookie.

Among the wealth of data gathered by *Inspire*, we have focussed on the accompanying text messages, extracted from the transcripts of 2557 negotiations. Each negotiation had two different participants, and one person participated in only one negotiation. The total number of contributors was over 5000; most of them were not native English speakers. The data contain 1,514,623 word tokens and 27,055 types. Compared with benchmark corpora, for example the Brown or the Wall Street Journal corpus (Francis and Kucera, 1997; Paul and Baker, 1992), this collection has a lower type-token ratio and a higher presence of content words among the most frequent words (this is typical of texts on a specific topic), and a high frequency of singular first- and second-person pronouns (this is typical of dialogues).

We considered all messages from one negotiation to be a single negotiation text. We concatenated the messages in chronological order, keeping the punctuation and spelling unedited. The communication thus captured was shaped by the fact that its messages were part of a negotiation process. Each nego-

Features	Split	NB				SVM				DT			
		Acc	F	P	R	Acc	F	P	R	Acc	F	P	R
negotiation-related	1/2 and 1/2	68.1	70.4	73.0	68.0	73.6	76.8	75.4	78.2	73.9	78.8	72.1	86.8
negotiation-related	3/4 and 1/4	69.1	71.3	74.1	68.7	73.7	77.0	75.5	78.5	75.4	79.4	73.8	86.0

Table 1: Classification accuracy and corresponding F-score, Precision and Recall. Classifying all negotiations as successful gives a baseline accuracy of 55%.

tiation was a training example in one of two classes: successful or unsuccessful. *Inspire* assigned a negotiation to the right class automatically. 55% of negotiations in our data set ended up with agreement (were successful).

We represented a complete negotiation, or text as we consider it, as a combined bag of words. We matched the tokens in the messages with an inventory of domains from Longman Dictionary of Contemporary English (Procter, 1978). This allowed us to select those terms that refer to negotiation specific issues – we call them *negotiation-related words*; we select *strategic words* based on words and patterns that literature shows to express the intentions, influence, self-obligations and motivations of the negotiation participants. In classifying successful and unsuccessful negotiations, subsets of these two types of features provided better accuracy than statistically selected features, e.g. most frequent unigrams and unigrams with a higher log-likelihood (Sokolova et al, 2005).

We halved each text, that is to say, the complete record of a negotiation. For each half we built a bag of 123 negotiation-related words – more on this in section 4. The binary attributes represented the presence or absence of the word in its half of the text. We concatenated the two bags, and labelled the resulting bag by the outcome of the whole negotiation: positive if the negotiation was successful, negative otherwise. We repeated this procedure for the split of the negotiation text into 3/4 and 1/4. Our ML tools were Weka’s (Witten and Frank, 2005) NAIVE BAYES (NB), the sequential minimal optimization (SVM) version of SUPPORT VECTOR MACHINE, and DECISION TREE (DT). In Table 1 we report the best results of the exhaustive search on adjustable parameters for every classifier; the evaluation method was tenfold cross-validation. Our accuracy results are comparable with those reported in previous studies (Kersten and Zhang, 2003; Nastase, 2006; Sokolova and Szpakowicz, 2006).

We used the paired *t-test* to generalize the results on both splits.¹ The two-tailed P value was 0.0102. By conventional criteria, this difference is considered to be statistically significant. Accuracy and, especially, Precision results show that DECISION TREE IS sensitive to the positions of words in different parts of the negotiations. The accuracy of SUPPORT VECTOR MACHINE and NAIVE BAYES changes only slightly. The precision and recall results give a better picture of the performance. The presence/absence of words recorded for different splits of negotiations influences the identification of true positive examples (successful negotiations) and true negative examples (unsuccessful negotiations). Successful negotiations have the highest rate of true identification achieved by DT, when the negotiations are split in half. Unsuccessful negotiations have the highest rate of true identification achieved by NB, when the split is 3/4 and 1/4; this split lets us improve the worst rates of true classifications – unsuccessful negotiations for DT and successful negotiations for NB. Generally, the unequal split allows us to reduce the difference between true positive and true negative results, and thus makes the classification of negotiations more balanced than the equal split.

4 The Empirical Set-up

We wanted to determine the placement of the segment of a negotiation most important in deciding whether the outcome is positive: at the beginning or at the end of the process. To do that, we split each negotiation in half, and built two parallel data sets, corresponding to the two halves. We classified each part using various ML tools. Next, we repeated the same classification tasks using smaller and smaller final segments, in order to monitor the variation in performance. Thus each negotiation text N consisted of the head segment (h) and the tail segment (t): $N = h \cup t, h \cap t = \emptyset$, where $|t| = \frac{|N|}{i}$ and t was

¹Results on the same data require the paired version of *t-test*.

the segment at the end of N , and $|h| = \frac{(i-1)|N|}{i}$ covering the beginning of the negotiation. We stopped when for two consecutive splits two classifiers had better accuracy on the head than on the tail. Each segment got the same class label as the whole negotiation.

For these experiments, as briefly explained in section 3, we took the textual negotiation data represented as bags of words. Because of the large number of word features (27,055 tokens), we performed lexical feature selection.

The statistical analysis of the corpus built from the *Inspire* negotiation messages has revealed that the issues discussed in these messages can be grouped into a small set of topics. The particular topic or domain to which a word belongs derives from the most frequent bigram and trigram meanings; for instance, the second most frequent trigram with the word *delivery* is *payment upon delivery*, so we assign *delivery* to the domain *negotiation process*. The data come from negotiations on a specific topic (sale/purchase of bicycle parts), so a likely candidate subset would be words related to it. The negotiation-related words constitute the first set of features. We show a text sample with the negotiation-related words in **bold**:

*Dear Georg, I hope you are doing well. I send you this message to ask you what happened to our **offer**. Just be aware that we will not be indefinitely waiting on your response. As I told you during our last **meeting**, IteX Manufacturing needs a partnership. So it is important to me to know if you are ready to **negotiate** with us. We can not **afford** losing so much precious time. We give you now five more **days** to answer our **offer** (1st of december 1997, 2400 Swiss time). After this dead line, we will propose our services to your concurrence. I still believe in a good partnership and relationship between our two societies. Let me know if you think so. For IteX Manufacturing. Rookie.*

Strategies which the negotiators adopt (promises, threats, exchange of information, argumentation, and so on) affect the outcome (Sokolova and Szpakowicz, 2006). Since the messages are dense, short and grammatically simple, the expression of strategies through language is straightforward and concentrates on communicating the main goal. The word categories that convey negotiators' strategies are modals, personal pronouns, volition verbs, mental verbs; we refer to them as *strategic* words. Strategic words constitute the second set of features. Our

text sample with strategic words in **bold** looks as follows:

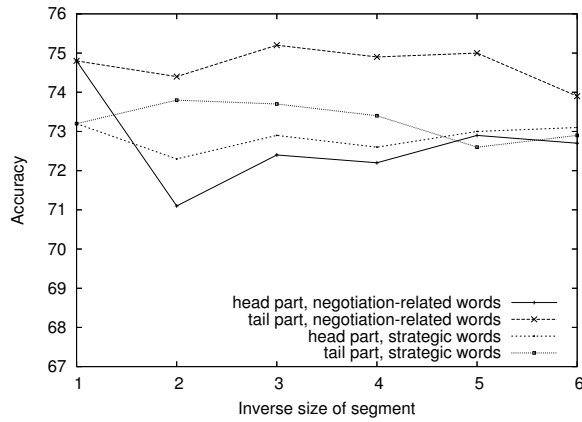
*Dear Georg, I hope **you** are doing well. I send **you** this message to ask **you** what happened to **our** offer. Just be aware that **we will not be** indefinitely waiting on **your** response. As **I** told **you** during **our last** meeting, IteX Manufacturing needs a partnership. So it is important to **me** to **know** if **you** are ready to negotiate with **us**. **We can not afford** losing so much precious time. **We** give **you** now five more days to answer **our** offer (1st of december 1997, 2400 Swiss time). After this dead line, **we will propose** our services to **your** concurrence. **I** still **believe** in a good partnership and relationship between **our** two societies. Let **me** know if **you** think so. For IteX Manufacturing. Rookie.*

We work with kernel (SVM), decision-based (DT) and probabilistic (NB) classifiers. Applying classifiers with different working paradigms allow us to capture and understand different aspects of the data, as the results and our discussion in section 5 will show. For each classifier, we used tenfold cross-validation and exhaustive search on adjustable parameters in model selection. The best results, in particular with high overall accuracy, appear in Figure 1.

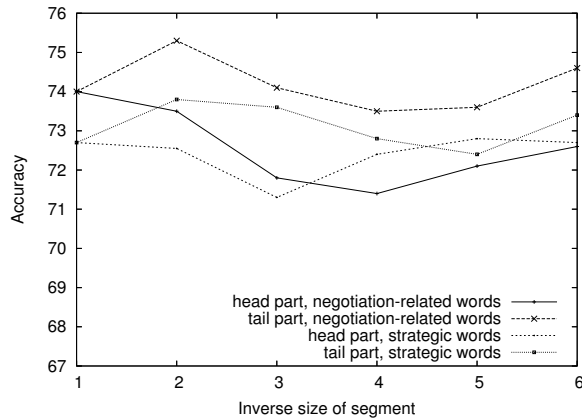
When the data are represented using negotiation-related words, the *tail* segments give more accurate outcome classification than the *head* segments. This holds for all splits and all classifiers; see Figure 1. The increase in accuracy when the head segments grow was to be expected, although it does not happen with DT and SVM – only with NB. At the same time, there is no monotonic decline in accuracy when the length of the tail segments decreases. On the contrary, NB constantly improves the accuracy of the classification. We note the fact that NB increases the accuracy on both head and tail segments and makes the basic assumption of the conditional independence of features. We explain the NB results by the decreased dependence between the presence/absence of negotiation-related words when the negotiations move to the second part of the process.

The results on the strategic-word representation are slightly different for the three classifiers; see Figure 1. SVM classifies all tail segments better than head segments, DT classifies tail segments better than head segments up to the $\frac{4}{5}/\frac{1}{5}$ split, and NB classifies the tail segment better than the head segment only for the half-and-half split. The accuracy results

1. DT



2. SVM



3. NB

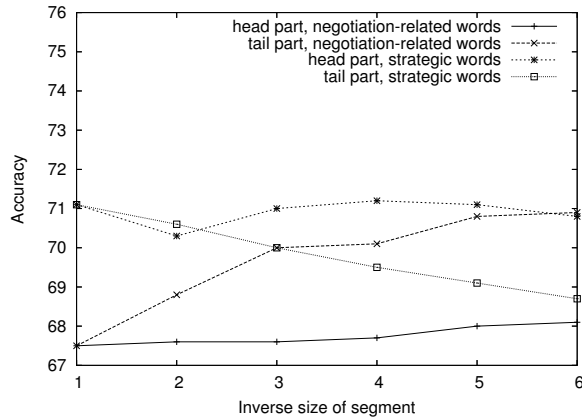


Figure 1: The classification accuracy with DT, SVM and NB, for negotiation-related and strategic words

are unstable for all three classifiers, with the accuracy on the head segments decreasing when the segments grow and the accuracy on the tail segments increasing when the tail segments shrink. The perfor-

Classifier	tail	s_1	s_2	s_1	s_2	s_3
DT	74.4	71.9	74.9	72.5	71.9	73.9
SVM	75.3	70.5	73.5	70.8	69.9	74.6
NB	68.8	68.5	70.1	68.7	68.9	70.9
Negotiation-related words						
Classifier	tail	s_1	s_2	s_1	s_2	s_3
DT	73.8	73.8	73.4	71.7	71.4	72.9
SVM	73.8	70.9	72.8	72.0	71.3	73.4
NB	60.8	70.6	69.5	69.2	69.3	68.7
Strategic words						

Table 2: The accuracy of the negotiation outcome classification on 2 and 3 splits of the second half of the negotiation – the *tail* segment. Classifying all negotiations as successful gives a baseline accuracy of 55%.

mance of the classifiers indicate that, as the deadline approaches, negotiation-related words reflect the negotiation process better than strategic words.

To investigate which part of the tail segments is more important for classifying the outcomes, we introduced additional splits in the tail segments. We divided the second half of each text into 2 and 3 parts and repeated the classification procedures for every new split. The results appear at the top of Table 2, where *tail* shows the classification results when the second half of the text was classified, and the other columns report the results on the tail splits; both splits satisfy the conditions $tail = \bigcup_i s_i$, where $s_i \cap s_j = \emptyset$ for every $i \neq j$.

The results show that adding splits in the *tail* segments emphasizes the importance of the last part of a negotiation: in both experimental settings, the classification of the outcome on the last part of the *tail* is more accurate than on its other parts. This holds for all three classifiers. For the strategic-word representation the same is true for SVM and partially for DT, but not for NB; see the bottom of Table 2. NB classifies the negotiation outcomes more accurately on splits s_1 than on splits s_2 and s_3 .

5 Segmentation Results

Taking into account the results reported in section 4, we chose negotiation-related words as the feature set. We selected for further analysis the half that performed better for a majority of the tools used. We focussed on the last part of the negotiation, and we extracted a gradually smaller fragment ($1/2 - 1/9$; 9 is the average number of text messages in one negotiation). Figure 2 plots the results of the experiments performed with decreasing segment sizes. As

Classifier	1/3		1/4		1/5		1/6		1/7		1/8		1/9	
	P	R	P	R	P	R	P	R	P	R	P	R	P	R
DT	74.2	85.3	74.2	84.3	75.2	82.3	73.61	83.02	74.53	82.39	72.1	81.62	74	81.3
SVM	76.1	78.1	76.3	76.3	77	75.3	78.3	75.3	77.2	73.4	76.9	72.3	77.6	71.6
NB	73.8	71.8	71.8	73.9	74.8	71.9	74.9	72	71.3	72.2	70.8	72.5	70.5	74.3

Table 3: Precision and recall on the “tail” segments; negotiation-related words.

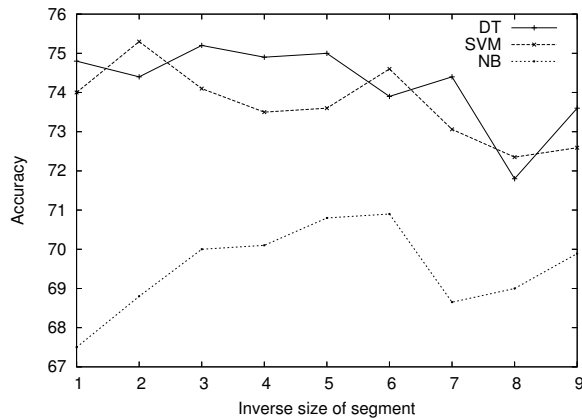


Figure 2: The evolution of accuracy with decreasing segment sizes.

we see, the *tail* segment of the length $1/7$ gave a decline of the accuracy for SVM and NB, with a slight improvement on smaller *tail* segments.

A more detailed analysis comes from looking at the precision and recall results on the segments; see Table 3. On $1/7$ and $1/9$ *tail* segments all classifiers have improved the identification of true negatives (unsuccessful negotiations). This means that the trends in the class of unsuccessful negotiations become more noticeable for the classifiers when the deadline approaches. The $1/8$ split is an exception, with the abrupt drop of true negative classification by DECISION TREE. The correct classification of positive examples (successful negotiations), however, diminishes when splits become smaller; this applies to the performance of all three classifiers. This means that at the end of the negotiations the class of successful negotiations becomes more diverse and, subsequently, multi-modal, and the trends are more difficult to capture by the classifiers.

As in the previous experiments, NB’s accuracy on the tail segments is higher than on the complete data. The opposite is true for SVM and DT – their accuracy on the *tail* segments is lower than on the complete data. We explain this by the fact that the sizes of

tail segments in the last splits do not give these two classifiers sufficient information.

6 Discussion and Future Work

We have analyzed textual messages exchanged in the course of electronic negotiations. The results support our hypothesis that texts of electronic negotiation have different characteristics than records of face-to-face negotiation. In particular, messages exchanged later in the process are more informative with regard to the negotiation outcome than messages exchanged at the beginning.

We represented textual records of negotiations by two types of word features. These features capture the important aspects of the negotiation process – negotiation-related concepts and indicators of the strategies employed. We performed extensive experiments with different types of ML algorithms and segments of varying sizes from the beginning and the end of the negotiation, on a collection of over 2500 electronic negotiations. Our study shows that words expressing negotiation-related concepts are more useful for distinguishing successful and failed negotiations, especially towards the end of negotiations. We also have shown that there is no linear dependency between the segment sizes and accuracy of the negotiation outcomes.

Our research plans include a continuation of the investigation of the negotiators’ behaviour in electronic negotiations and its reflection in language. To see whether dialogue analysis improves prediction of the negotiation outcomes, we will look at negotiations as dialogues between participants and take into account their roles, e.g. buyer and seller. We will split the negotiation at message boundaries to avoid arbitrary split-points of the negotiation process.

Acknowledgments

Partial support for this work came from the Natural Sciences and Engineering Research Council of Canada.

References

- W. Adair, J. M. Brett. 2005. The negotiation dance: time, culture, and behavioral sequences in negotiation. *Organization Science*. 16(1): 33-51.
- J. Blaták, E. Mráková, L. Popelínský. 2004. Fragments and Text Categorization. *Proceedings of the 42th Annual Meeting of the Association of Computational Linguistics (ACL 2004)*. Companion Volume: 226-229, Association for Computational Linguistics.
- J. M. Brett. 2001. *Negotiating Globally*. Jossey-Bass.
- S. R. Charvet. 1997. *Words that Change Minds: Mastering the Language of Influence*. Kendall/Hunt.
- J. Chu-Carroll, S. Carberry. 2000. Conflict Resolution in Collaborative Planning Dialogues. *International Journal of Human-Computer Studies*. 53(6): 969-1015.
- W. N. Francis, H. Kučera. 1967. *Computational Analysis of Present-Day American English*, Brown University Press.
- G. E. Kersten, G. Zhang. 2003. Mining Inspire Data for the Determinants of Successful Internet Negotiations. *Central European Journal of Operational Research*. 11(3): 297-316.
- S. Koeszegi, E.-M. Pesendorfer, R. Vetschera. 2007. Data-driven Episodic Phase Analysis of E-negotiation. *Group Decision and Negotiation 2007*. 2: 11-130.
- V. Nastase. 2006. Concession curve analysis for Inspire negotiations. *Group Decision and Negotiation*. 15: 18-193.
- D. B. Paul and J. M. Baker. 1992. The Design for the Wall Street Journal-based CSR Corpus. *Proceedings of the 2nd International Conference on Spoken Language Processing (ICSLP'92)*, 357-361.
- P. Procter. 1978. *Longman Dictionary of Contemporary English*. Longman Group Ltd. Essex, UK.
- D. Reitter, J. Moore. 2007. Predicting Success in Dialogue. *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics (ACL 2007)*, 808-815, Association for Computational Linguistics.
- T. Simons. 1993. Speech Patterns and the Concept of Utility in Cognitive Maps: the Case of Integrative Bargaining. *Academy of Management Journal*. 38(1): 139-156.
- M. Sokolova, V. Nastase, M. Shah, S. Szpakowicz. 2005. Feature Selection in Electronic Negotiation Texts. *Proceedings of Recent Advances in Natural Language Processing (RANLP'2005)*, 518-524, Incoma Ltd, Bulgaria.
- M. Sokolova, S. Szpakowicz. 2006. Language Patterns in the Learning of Strategies from Negotiation Texts. *Proceedings of the 19th Canadian Conference on Artificial Intelligence (AI'2006)*, 288-299, Springer.
- M. Sokolova and S. Szpakowicz. 2005. Analysis and Classification of Strategies in Electronic Negotiations. *Proceedings of the 18th Canadian Conference on Artificial Intelligence (AI'2005)*, 145-157, Springer.
- , H. Takeuchi, L. Subramaniam, T. Nasukawa, S. Roy. 2007. Automatic Identification of Important Segments and Expressions for Mining of Business-Oriented Conversations at Contact Centers. *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, 458-467, Association for Computational Linguistics.
- I. Witten, E. Frank. 2005. *Data Mining, 2nd ed.*. Morgan Kaufmann. www.cs.waikato.ac.nz/ml/weka/
- L. Young. 1991. *Language as Behaviour, Language as Code: A Study of Academic English*. John Benjamins.