

Negotiation Dynamics: Analysis, Concession Tactics, and Outcomes

Koen Hindriks, Catholijn M. Jonker, Dmytro Tykhonov
Delft University of Technology, Man-Machine Interaction
{k.v.hindriks, c.m.jonker, d.tykhonov}@tudelft.nl

Abstract

Given that a negotiation outcome is determined to a large extent by the successive offers exchanged by negotiating agents, it is useful to analyze dynamic patterns of the bidding, what Raiffa calls the “negotiation dance”. Patterns in such exchanges may provide additional insight into the strategies used by the agents. The current practice of evaluating a negotiation strategy, however, is to primarily focus on fairness and quality aspects of the agreement. There is a lack of tools and methods that facilitate a precise analysis of the negotiation dynamics. To fill this gap, this paper introduces a method for analysis based on a classification of negotiation steps.

The method provides the basic tools to perform a detailed and quantified analysis of a negotiation between two agents in terms of dynamic properties of the negotiation trace. The method can be applied to well-designed tournaments, but can also be used to analyze single 1-on-1 negotiation.

Example findings of applying the method to analyze the ABMP and Trade-Off strategies show that sensitivity to the preferences of the opponent is independent, respectively dependent, on a correct model of that opponent. Furthermore, the results illustrate that having domain knowledge is not always enough to avoid making unintentional steps.

1. Introduction

The negotiation dance of exchanging successive offers by negotiation partners affects the negotiation outcome [11]. To gain more insight in the negotiation dynamics, in [1] a classification of negotiation steps was introduced in order to characterize and compare the bidding process of humans and software agents. The results showed an overall similarity of the bidding style of humans and the Agent-Based Market Places (ABMP) strategy, a concession- oriented negotiation strategy, see [7]. However, the analysis did not provide insights in why the different kinds of steps were made, nor did it help us understand why and to what extent these steps affect the outcome of the negotiation. As far

as we know, no analytical methods exist that do provide the desired insights.

The step-wise analysis method introduced in this paper is a concrete step towards providing such insights. It extends the work presented in [1], by extending and providing a precise characterization of the negotiation step classification and by providing some useful metrics. These metrics in turn are used to define more complex dynamic properties of the negotiation dance to facilitate the analysis of various dynamic properties of the strategies under evaluation.

Other analytical methods mentioned in the literature typically assess the performance of negotiation strategies in terms of fairness and quality aspects of the agreement (if any) that agents reach. Aspects considered are who wins, the distance of the outcome to the Pareto Efficient Frontier, the Nash Product, and the Kalai-Smorodinsky Point (see section 4.3 for details). Formal definitions of these concepts can be found in e.g., [11]. Such measures of evaluation focus on the *negotiation outcome*.

Instead, the concepts introduced here are intended to facilitate the analysis of typical bidding patterns induced by various negotiation strategies. It is the objective of this paper to propose a method and some metrics that facilitate a *precise characterization* of the negotiation dance. In turn, such a characterization of the dynamics of negotiation may contribute to the identification of explanations for such findings. It is the aim of this paper to at least *partially* identify some of the *reasons* that may explain particular findings, that is, to associate particular aspects of a negotiation problem or strategy with particular extreme values (e.g., minimum or maximum) of the metrics defined below.

We illustrate the use of these concepts for the analysis of *concession tactics*. For example, although it is generally acknowledged that a concession should actually *increase* the utility of the opponent and not just be a move that *decreases* one’s own utility, in practice, as we will show, such behavior is not always achieved by strategies that have been designed to concede towards the opponent. Steps that reduce both the agent’s own as well as its opponent’s utility have been called *unfortunate steps* (cf. [1]). Both humans as

well as software agents using the ABMP strategy were observed to make such steps in negotiation experiments reported in [1], but humans made fewer of them. From this analysis we conclude that it is impossible to avoid unfortunate steps without sufficient domain knowledge or a model of the negotiation partner.

The paper is organized as follows. The next section discusses related work. In Section 3, we briefly introduce the topic of negotiation dynamics. Section 4 introduces the step-wise analysis method and some metrics for analyzing dynamic negotiation properties. In Section 5, the method is illustrated by analyzing the Trade-Off [5] and ABMP [7] strategy in various negotiation domains. Finally, the paper concludes with some suggestions for research on automated negotiation derived from the proposed analysis method.

2. Related Work

In order to develop efficient negotiation strategies that are robust as well as outperform other strategies, we believe it is important to be able to evaluate the dynamic behaviour induced by negotiation strategies. An analysis of negotiation dynamics in our opinion requires the use of both theoretical as well as experimental evaluation methods, in which at least the following aspects are attended to:

- *competition* with other strategies and itself,
- *case studies* of varying complexity, and
- *theoretical properties* of the dynamics.

This paper discusses some case studies of varying complexity and some theoretical properties. For papers that focus on the competition with other strategies the reader is referred to e.g., [5, 6, 14]. Additionally, the scope of research reported here concerns the negotiation dynamics as a pattern of offers (cf. [11]). That is, our work concerns *bargaining*, a method for reaching joint agreements by means of exchanging offers according to e.g., an alternating offers protocol.

There are many papers discussing the quality of negotiation strategies. We briefly mention some papers that discuss aspects that are complementary to those discussed here. In [9], a classification scheme is provided that defines some properties that are oriented towards *rationality and the use of resources*. In [8], the authors discuss the performance and evaluation of a negotiation software agent against humans. The evaluation measures used concern the agent's utility of the agreement versus that of humans, and the *social welfare* associated with the agreement.

In [1], a formalization of the negotiation process is provided together with a set of performance properties that facilitate evaluation of the quality of the agreement

reached, based on the work of [12, 14]. The paper also discusses some dynamic properties of the bidding. Similar to the work presented in [8], the authors used the SAMIN system, inspired by the system presented in [6], to analyze the ABMP strategy playing against itself and playing against human negotiators. The experiments showed that human and ABMP negotiators primarily made *concession steps* (see for a precise definition Section 4). Additionally, it was shown that humans were more diverse, i.e. the types of negotiation steps they performed were more diversified.

The initial, informal classification of negotiation steps and the results reported in [1] form the inspiration of the current paper. In combination with the belief expressed by many that the pattern of offers exchanged influences the negotiated outcome, this motivated our study of negotiation strategies from the perspective of the negotiation dynamics and the actual steps made.

3. Negotiation Dynamics

In the analysis of negotiation strategies, not only the outcome of a negotiation is relevant, but also the bidding process itself is important. Mistakes made during the bidding can have an enormous impact on both players. Although experienced negotiators confirm this, and it is also recognized by researchers of negotiation strategies for automated negotiation, this hypothesis is difficult to quantify. Examples from human negotiations are of the form: "a wrong offer can upset relationships, even causing the other party to walk away", or "Sometimes an offer that is meant as a concession to the other party confuses the issues. This can only be circumvented if there is enough trust between the parties to exchange some information on their respective preferences."

From the point of view of automated negotiation, the objective is to stay as close as possible to the Pareto Efficient Frontier. However, in automated negotiations typically no prior information is exchanged about the preferences of the negotiating parties, and none of the players know where the Pareto Efficient Frontier actually is. It thus remains a challenge to stay or end close to that Frontier. To this end, opponent modelling is an often used approach to better predict which bids will be appreciated by the other party, see, e.g., [3, 10].

More precisely, five key factors can be identified that shape the outcome of a bilateral negotiation with incomplete information: (i) knowledge about the negotiation domain (e.g. the market value of a product or service), (ii) one's own and one's opponent's preferences, (iii) process attributes (e.g. deadlines), (iv) the negotiation strategies, and (v) the negotiation protocol. In this paper, our interest is in analyzing,

classifying and in precisely characterizing aspects of the negotiation dynamics that influence the final agreement of a negotiation. The main interest thus is in proposing concepts and metrics that relate these factors to specific aspects of the negotiation dynamics and to thus gain a better understanding of the final outcome of a negotiation. The analysis does not take the features of the protocol per se into account but instead focuses on the exchange of offers. In principle, the method allows for generalizations to multi-party negotiations but we do not consider such extensions here (see e.g. [13]).

4. Step-Wise Analysis Method

In bilateral bargaining, the negotiation dynamics is completely represented by the sequence of offers $t = \langle b^1_s, b^2_o, b^3_s \dots \rangle$ exchanged between parties S and O , also called the *negotiation trace*. A negotiation trace is called *closed* if it ends in either an *accept* or *withdraw* move by either party. In this section the basic notions of the step-wise analysis method are defined: classes of negotiation steps, metrics, outcome properties, and patterns over negotiation steps. After that the different phases of the method are defined.

4.1 Negotiation steps

The key unit of negotiation dynamics analysis is a *single negotiation step* performed by one of the negotiating parties. A negotiation step in bargaining consists of an offer proposed by one party to the other. If this offer is not the first offer proposed by an agent, it typically is computed using at least the previous offer proposed by that agent as input. To record this fact and to facilitate notation below, formally, a negotiation step s by agent a is modeled as a *transition* from a previous offer b_a to a newly proposed offer b'_a , which is written as $b_a \rightarrow b'_a$. Such steps can be classified based on the associated utility for both parties.

For the step-wise analysis method the classes of negotiation steps of [1] have been extended with two additional classes. In a number of papers, it has been suggested, that it is smart not to make concessions too soon, but to move over one's iso-utility lines first [5]. For this reason, a separate category of *nice steps* that move in the direction of the opponent but do not concede own utility is introduced. An example of a strategy that is designed to make such steps is the Trade-Off Strategy based on similarity criteria discussed in [5], a variation is proposed in [13]. Additionally, so-called *silent steps* are introduced to represent the fact that parties sometimes repeat their offers, and may not make any concessions at all such as in a Boulware or *take it or leave it* strategy, see [12].

Before formally defining the concepts below, some additional notation is introduced. $U_S(b)$ denotes the utility of "Self" with respect to bid b . Similarly, $U_O(b)$ denotes the utility of "Other" with respect to b . We use $\Delta_a(b, b') = U_a(b') - U_a(b)$, $a \in \{S, O\}$, to denote the utility difference of two bids b and b' in the utility space of agent a . We also write $\Delta_a(s)$ to denote $\Delta_a(b, b')$ for a step $s = b \rightarrow b'$. Here we present a precise definition of the classes of negotiation steps proposed in [1] extended as discussed above. These step categories define the core of the step-wise analysis method.

Definition of Step Classes:

Let $s = b_s \rightarrow b'_s$ be a step in the bidding by Self (the definition for Other is completely symmetric). Then the negotiation step s taken by Self is classified as a:

- *Fortunate Step*, denoted by (S+, O+), iff:
 $\Delta_S(s) > 0$, and $\Delta_O(s) > 0$.
- *Selfish Step*, denoted by (S+, O_≤), iff:
 $\Delta_S(s) > 0$, and $\Delta_O(s) \leq 0$.
- *Concession Step*, denoted by (S-, O_≥), iff:
 $\Delta_S(s) < 0$, and $\Delta_O(s) \geq 0$.
- *Unfortunate Step*, denoted by (S_≤, O-), iff:
 $\Delta_S(s) \leq 0$, and $\Delta_O(s) < 0$.
- *Nice Step*, denoted by (S=, O+), iff:
 $\Delta_S(s) = 0$, and $\Delta_O(s) > 0$.
- *Silent Step*, denoted by (S=, O=), iff:
 $\Delta_S(s) = 0$, and $\Delta_O(s) = 0$.

The proposed classification is exhaustive, and all step classes are disjoint. (To allow for some marginal errors the areas of the Nice and Silent steps can be stretched somewhat in the analysis. In that way, a step in which only 0.005 of Self's utility is lost would still be classified as e.g. Nice, instead as a Concession.) In a concession step some own utility needs to be conceded but the opponent's utility may stay the same. In such cases, Self can claim that it made a concession step by arguing that it conceded some of its own resources.

Each step type in a negotiation typically has a distinct role or function, though often in automated negotiation systems not all of these step types are taken into account. Fortunate steps happen spontaneously in human negotiations (see [1]). Having a strategy that is able to perform such steps deliberately is beneficial, since such steps can be used to recover from moves away from the Pareto Efficient Frontier, e.g., as the result of *concessions* or *unfortunate steps*. The latter two steps aim at reaching a jointly acceptable outcome. Although it would be best to avoid unfortunate steps at all times when conceding, it is impossible to guarantee this when the opponent's preferences are known incompletely. Selfish steps may be performed by an

agent to signal to the other party that a previous move is not appreciated. The role of nice and silent steps has been discussed above.

4.2 Step Metrics and Pattern Properties

Having established different types of negotiation steps that are useful in the analysis of negotiation strategies, we now introduce and define metrics in terms of these steps that can be used for the analysis of negotiation traces. First, some additional notation is defined. Given a trace $t = \langle b^1_s, b^2_o, b^3_s \dots \rangle$ of offers, t^i denotes the i^{th} element of this sequence. Let t_s (resp. t_o) denote the sequence of steps from t that are made by agent ‘‘Self’’ (resp. ‘‘Other’’) and let class $c \in \{\text{Fortunate, Nice, Concession, Selfish, Unfortunate, Silent}\}$; then t_c denotes the subsequence of steps that belong to class c . Finally, $t_{(a,c)}$, also written t_{ac} , denotes the subsequence of steps by $a \in \{S, O\}$ that belong to class c . The following step metrics are introduced here:

Definition. Number of Steps per Trace

The number of steps $\#t$ in a trace t of length $|t| = n$ is defined as follows: $\#t = |t| - 1$.

Definition. Total Utility per Class

The pair $Total_c(t)$ of sums of utility differences in all steps of class c in a sequence t of steps is defined by:

$$Total_c(t) = (Total_{sc}(t), Total_{oc}(t)),$$

where for any agent $a \in \{S, O\}$: $Total_{ac}(t) = \sum_i \Delta_a(t^i_c)$.

Definition. u-Average Utility per Class

The pair $u-Ave_c(t)$ of average differences in utility in all steps in class c in a sequence t of steps is defined by:

$$u-Ave_c(t) = (u-Ave_{sc}(t), u-Ave_{oc}(t)),$$

where for any agent $a \in \{S, O\}$:

$$u-Ave_{ac}(t) = \sum_i \Delta_a(t^i_c) / \#t_c.$$

Here $\#t_c$ is a number of steps of class c in trace t . This metric measures the average utility conceded per negotiation step. A relative measure can be defined in terms of this metric to identify how much utility has been conceded by agent a relative to the other, indicated by agent g : $u-Ave_c(t_a) / u-Ave_c(t_g)$. This figure, if not identical to 1, indicates that one party is a *conceder* relative to the other, and that concessions may not have been paced and linked to that of the other party, as is advised by Raiffa [11, p. 128].

Definition. % per Class

The percentage $\%_c(t)$ of class c steps in a trace t is defined by: $\%_c(t) = \#t_c / \#t$.

Negotiation strategies can be designed with specific aims in mind that should be observable as patterns in the negotiation dance. For example, the success of a

strategy that is supposed to learn its opponent’s preferences can be verified by checking whether the frequency and/or size of unfortunate steps over a negotiation trace decreases. Such patterns can be seen as a measure of *adaptability of a party to its opponent*. Another useful measure of the *sensitivity to the opponent’s preferences* can be defined by comparing the percentage of *fortunate, nice and concession steps* that increase the opponent’s utility to the percentage of *selfish, unfortunate and silent steps* that decrease it. Intuitively, an agent that only performs steps that increase its opponent’s utility can be said to be (very) sensitive to the needs of its opponent.

Definition. Sensitivity to Opponent Preferences

The measure for sensitivity of agent a to its opponent’s preferences is defined for a given trace t by:

$$sensitivity_a(t) = \frac{\%_{Fortunate}(t_a) + \%_{Nice}(t_a) + \%_{Concession}(t_a)}{\%_{Selfish}(t_a) + \%_{Unfortunate}(t_a) + \%_{Silent}(t_a)}$$

In case no selfish, unfortunate or silent steps are made we stipulate that $sensitivity_a(t) = \infty$. If $sensitivity_a(t) < 1$, then an agent is more or less insensitive to opponent preferences; if $sensitivity_a(t) > 1$, then an agent is more or less sensitive to the opponent’s preferences, with complete sensitivity for $sensitivity_a(t) = \infty$. Typically, this sensitivity measure varies with different domains and different opponents and averages over more than one trace need to be computed. Note that the notion of sensitivity is *asymmetric*: one agent may be sensitive to the other’s preferences, but not vice-versa. In section 5, this metric is used to analyze the sensitivity of two existing negotiation strategies and its relation to knowledge of the opponent’s preferences and the negotiation outcome is discussed.

4.3 Outcome properties

The usual outcome performance metrics are used to determine the quality of an agreement with respect to players, the distance to the Pareto Efficient Frontier, Nash Product, and Kalai-Smorodinsky, see e.g., [9, 2]. Other global measures taken are the (average) number of negotiation rounds R needed to reach an agreement, the number of agreements A reached in a tournament, and the time T taken by each party. Outcome properties are not the main topic of this paper.

4.4 Phases of the Step-Wise Analysis Method

Based on the above concepts, the step-wise analysis method is specified by:

1. Starting point: the strategies to be analysed, and a library of domains and of other strategies that can be used to test the input strategies.

2. If necessary, implement the input strategies.
3. Set up a tournament with the selected negotiation strategies and case studies.
4. Run the tournament and log every negotiation, along with the names of the strategies, the domain name, and the preferences of the players.
5. Test and classify the step properties and metrics, classifying all steps of a negotiation and logging the results, test the outcome properties and the other overall properties
6. Apply statistics and produce analytical results.
7. Interpret results, produce graphics.

The effectiveness of the step-wise method is illustrated in the next section by some illustrative case studies. A tournament a la [6] is used to experiment with various strategies. In the tournament, strategies play against each other, against themselves, and are applied to varying negotiation domains, with varying preference profiles. The negotiation problems used are varied also in the amount of *domain knowledge* that is available. For example, the issue price typically is rather predictable, more is better for the seller, less is better for the buyer, and the normal ordering of the real numbers is maintained. For some domains one also knows that some issues are more important than others. This kind of knowledge is important in strategies such as the Trade-Off strategy [5], but other strategies do not employ this kind of knowledge.

5. Step-wise Analyses Method Applied

This section illustrates the method for a combination of strategies and negotiation domains, while focusing on individual steps – in particular *unfortunate steps* – and on the sensitivity of two negotiation strategies with respect to preferences of the opponent.

5.1 Strategies and Domains Used

The negotiation domains used in this paper are:

The *Second hand car selling* domain, taken from [7], includes 5 issues. Only the buyer’s preferences and the price issue are predictable, in the sense that an agent can reliably predict the other agent’s preferences associated with an issue.

The *Service-Oriented Negotiation* domain, taken from [5], includes 4 issues. All issues are predictable, i.e. based on available “domain knowledge” preferences can be reliably predicated.

The *AMPO vs City* domain, taken from [12], including 10 issues, of which only 8 are predictable. Information about the opponent’s issue priorities is not available, i.e. the weights agents associate with issues.

The following strategies have been studied: The *ABMP* strategy [7], a concession oriented strategy, which computes bids to offer next without taking domain or opponent knowledge into account. (Experiments were run with a negotiation speed of 0.1 and a concession factor of 1, see [7].)

The *Trade-off* strategy is based on similarity criteria [5], and exploits domain knowledge to stay close to the Pareto Frontier. The “smart” version of this strategy performs nice steps if possible; otherwise it concedes a fixed amount 0.05 (cf. [5]). For the Service-Oriented Negotiation domain, we reproduced the results presented in [5].

The *Random Walker* strategy randomly jumps through the negotiation space, and can be run with or without a break-off point (to avoid making offers below that utility). Random Walker serves as a “baseline” strategy.

5.2 Experimental results and interpretations

A tournament with the strategies and domains of the previous section was set up and run. A full analysis was made of the type of steps made, which was then used to calculate the average sensitivity ratio for all tested strategies over multiple runs against all strategies (including itself) in the domains described above. In that setting, ABMP shows an overall sensitivity of 2.5, Trade-off 2.2, and Random Walker 1.1. Theoretically, over all domains and against all strategies, Random Walker would have a sensitivity value of 1. When considering overall utility the scores are: ABMP 0.72, Trade-Off 0.74, and Random Walker 0.69. Note, that the efficiency of the agreement does not correlate strongly with the sensitivity values of the strategies. The sensitivity scores of ABMP and Trade-Off can be better understood by considering the domains in which they played. Figure 1 shows typical runs in the AMPOvsCity domain.

Figure 1a shows a run of Trade-Off, representing the City, versus Random Walker (with break-off set to 0.6), playing AMPO. The Random Walker strategy is insensitive with respect to its own preferences. This fact, combined with the lack of information of relative importance of issues (weights) causes the unfortunate steps produced by the Trade-Off strategy.

Figure 1b shows Trade-Off (as City) vs ABMP (as AMPO) in which ABMP is rather insensitive (average value 1.5), and Trade-Off is sensitive (average value 2.7). In this domain Trade-Off really exploits the available domain knowledge.

Figure 1c shows Random Walker (City) vs ABMP (AMPO). ABMP always concedes on all issues, determining the size of the concession on the difference

between the utilities of its own bid and that of its opponent. It does not use previous opponent bids to get insight into the opponent's preferences and as a result does not adapt much to the strategy of the opponent.

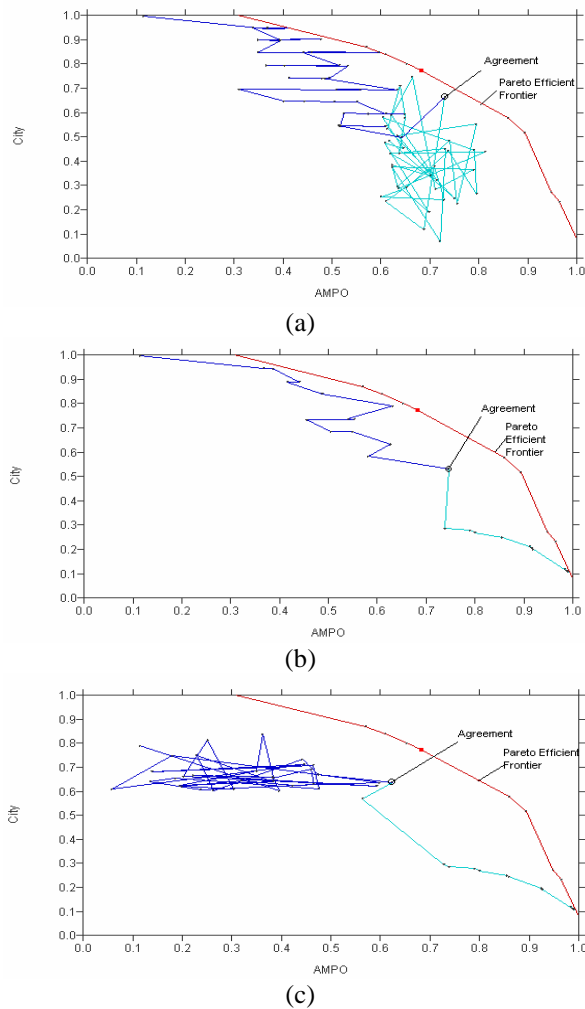


Figure 1 – Dynamics of negotiation process for: a) Trade-Off (City) vs Random Walker strategy (AMPO), b) Trade-Off (City) vs ABMP strategy (AMPO), c) Random Walker (City) vs ABMP strategy (AMPO). The Pareto Efficient Frontier is built according to its definition (see [12], pp. 227) using exhaustive search.

The opponent sensitivity analysis shows a direct link between the correctness and/or completeness of the domain knowledge and opponent preferences sensitivity. The Trade-Off strategy is very sensitive to opponent preferences given complete information. In that case, the similarity functions exactly match the opponent's preferences and the weights exactly represent the issue importance factors of the opponent and the sensitivity metric is converging to infinity.

The Service oriented negotiation domain does not have information about weights of the similarity

functions and thus opponent preferences sensitivity of the Trade-Off strategy decreases (2.7 in average) but it is still more sensitive to the opponent preferences than ABMP (1.5 in average). Similarity functions for the Second hand car domain were defined in such way that they often do not match the preferences of the negotiation opponents. In addition, the weights of the similarity function do not match the opponent's importance factors of the negotiation issues. This leads to underperformance of the Trade-Off strategy (1.5 in average) while ABMP shows more robust negotiation behavior (2.3 in average). The experiments show that if less domain knowledge is available, Trade-Off makes more unfortunate steps.

In general, when issues are predictable, the chance of making an unfortunate step becomes small. This aspect becomes clear in the car domain, where the seller's preferences are rather predictable, but the buyer's preferences vary a lot. ABMP's average sensitivity in the car domain to opponent's preferences is 2.3, but when restricting that to AMBP playing the buyer we get an average of 4.3.

We conclude that it is impossible to avoid unfortunate steps without sufficient domain knowledge or opponent knowledge. Indeed, the similarity criteria functions used in the Trade-Off Strategy provide general information about the negotiation problem, but do not take into account the specific attributes of the negotiating parties. In any particular case, a negotiator may deviate from the generalized domain model in various ways. Approaches as reported in [3, 8, 16] apply techniques to learn more about the opponent.

6. Conclusion

This paper shows that an analysis of the negotiation dance [12] is important for the understanding and improvement of negotiation strategies. The step-wise analysis method introduced in this paper focuses on the classification of negotiation steps and a metrics over this classification. The classification enables us to relate the *intent* of a strategy in making a negotiation step with the actuality of the *perception* of that step by the opponent. For example, a strategy might be concession oriented, i.e., steps are intended to be concessions, but in reality some of these steps might be *unfortunate*, meaning that although the proposer of the bid is giving in, from the perception of the receiver, the bid is actually worse than the previous bid.

By testing strategies over various domains and against various opponents patterns emerge of when such unfortunate steps occur. Furthermore, it gives more insights into properties such as the sensitivity of strategies. Experiments with the step-wise method

show, for example, that the Trade-Off strategy is rather *responsive* to the behavior of the opponent, in that it follows the behavior of the opponent. If that is rather wild, such as the random behavior of the Random Walker, it results (for both parties) in high percentages of unfortunate steps.

Experiments further show that the occurrence of unfortunate steps is related to features of the negotiation domain and the extent to which such features are incorporated in the strategy. The same holds for knowledge about the preferences of the opponent. To better understand the relative importance of each relation, we have emphasized the distinction between *domain knowledge* and *opponent knowledge*.

We think it is impossible to avoid unfortunate steps without sufficient domain knowledge or opponent knowledge. Domain knowledge provides generalized information about the negotiation problem, but does not necessarily match with individual preferences of negotiating parties. Opponent knowledge concerns individual information and as such is not transferable to other opponents. Therefore, we advocate a combination of domain and opponent knowledge.

The step-wise analysis method focusses explicitly on properties of interest to the researcher. The combination of statistical methods and graphical representation is strong: Humans process graphs faster than tables with numbers, however, the number of experiments typically done make it impossible for the human to view every graph produced by the experiment. The example in the paper shows how the focus on the percentage of unfortunate steps makes it possible for the step-wise analysis method to present to us the graph of the Trade-Off vs Random Walker, providing insight into the sensitivity of the Trade-Off strategy. A more general aspect is that graphs of long negotiation dances become hard to grasp; what may look like a neat series of nice steps, might actually be a mixture of unfortunate and nice steps.

We believe that our results also show the need for *benchmark problems* for bilateral negotiation.

Future work It would be interesting to experiment with the approach of [3] that extends the Trade-Off strategy with Kernel Density Estimates to learn the preference weights of the opponent. This approach and other strategies that incorporate learning (e.g. [8, 15]) should show a decrease over time of the frequency of unfortunate steps.

7. References

[1] T. Bosse, and C.M. Jonker, "Human vs. Computer Behaviour in Multi-Issue Negotiation", In: *Proceedings of the 1st International Workshop on Rational, Robust,*

and Secure Negotiations in Multi-Agent Systems, IEEE Computer Society Press, 2005, pp. 11-24.

[2] T. Bosse, C.M. Jonker, and J. Treur, "Experiments in Human Multi-Issue Negotiation: Analysis and Support", in: *Proc. of the 3rd Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, IEEE Computer Society Press, 2004, pp. 672-679.

[3] R.M. Coehoorn, and N.R. Jennings, "Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs", in: *Proceedings of the 6th International Conference on E-Commerce*, Delft, The Netherlands, 2004, pp. 59-68.

[4] P. Faratin, C. Sierra, and N. R. Jennings, Negotiation Decision Functions for Autonomous Agents, in: *Int. Journal of Robotics and Autonomous Systems*, 24 (3-4) 159-182, 1998.

[5] P. Faratin, C. Sierra, and N. Jennings, "Using Similarity Criteria to Make Negotiation Trade-Offs", *Journal of Artificial Intelligence*, 142 (2), 2003, pp. 205-237.

[6] P. Henderson, S. Crouch, R.J. Walters, and Q. Ni, "Comparison of some Negotiation Algorithms using a Tournament-Based Approach", in: *Agent Technologies, Infrastructure, Tools and Applications for E-Services*, Springer-Verlag, 2003, pp. 137 – 150.

[7] C.M. Jonker and J. Treur, "An Agent Architecture for Multi-Attribute Negotiation", In: *Proc. of the 17th Int. Joint Conference on AI (IJCAI'01)*, 2001, pp. 1195-1201.

[8] R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry, "An Automated Agent for Bilateral Negotiation with Bounded Rational Agents with Incomplete Information", in: *Proc. of the 17th European Conference on Artificial Intelligence (ECAI'06)*, 2006, pp. 270-274.

[9] A.R. Lomuscio, M. Wooldridge, and N.R. Jennings, "A classification scheme for negotiation in electronic commerce", in: *Int. Journal of Group Decision and Negotiation*, 12(1), 2000, pp. 31 – 56.

[10] C. Mudgal and J. Vassileva, "Bilateral Negotiation with Incomplete and Uncertain Information", in: *Proc. of the Workshop on Cooperative Information Agents*, Springer, 2000, pp. 107-118.

[11] H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press, 1982.

[12] H. Raiffa, J. Richardson, and D. Metcalfe, *Negotiation Analysis: The Science and Art of Collaborative Decision Making*, Harvard University Press, 2002.

[13] R. Ros and C. Sierra, "A Negotiation Meta Strategy Combining Trade-off and Concession Moves", *Journal of Autonomous Agent and Multiagent Systems*, 12(2), 2006, pp. 163-181.

[14] J.S. Rosenschein, and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*, MIT Press, 1994.

[15] D. Zeng and K. Sycara, "Bayesian Learning in Negotiation", in: *Proc. AAI Symposium on Adaptation, Co-evolution, and Learning in Multiagent Systems*, AAAI Press, 1996.

[16] L.L. Thompson, *The Heart and Mind of the Negotiator*, Pearson Prentice Hall, 2005.