

Computing negotiation update semantics in multi-issue bargaining dialogues

Volha Petukhova¹, Harry Bunt² and Andrei Malchanau¹

¹Spoken Language Systems Group, Saarland University, Germany

²Tilburg Center for Communication and Cognition, Tilburg University, The Netherlands

{v.petukhova, andrei.malchanau}@lsv.uni-saarland.de
harry.bunt@uvt.nl

Abstract

This paper presents a computational approach to modelling pragmatic and semantic aspects of multi-issue bargaining dialogues. The model accounts for actions that shape negotiation structure and actions that express negotiation strategies. The model also accepts a number of negotiation moves as specifications of the semantic content of the performed task-related dialogue acts. The designed dialogue context model specifies the creation, maintenance and transfer of participants' private and shared beliefs. A negotiation agent that operates on this basis was implemented and evaluated against human performance. The approach allows efficient interpretation and generation of negotiation behaviour according to different negotiation strategies.

1 Introduction

The fundamentals of human dialogue modelling are concerned primarily with the modelling of conversational goals and intentions, dialogue structure, grounding mechanisms and reasoning with the assumptions of rationality and cooperation. Dialogue models are important for interactive human-computer systems development. Most research in human-computer interaction modelling and dialogue systems design so far has been done in the area of task-oriented systems (TOS) with well-defined tasks in restricted domains.

Research efforts in dialogue modelling have recently been moving towards a world of smart environments seeking new ways of interfacing and engaging with technologies that more closely reflect rich natural human interaction in domains and settings of various complexity. The research community is targeting more flexible adaptable open-domain dialogue modelling driven by cognitive

modelling of human dialogue behaviour. Existing two-party TOS dialogue models are undergoing changes to reflect advanced understanding and to allow efficient computation of phenomena specific to new domains, new ways of interacting, and novel user experiences. For instance, it has been acknowledged that the assumption that conversational agents act fully rationally and cooperatively does not hold in many conversational settings, see e.g. (Traum et al., 2008b) and (Asher and Quinley, 2011). In competitive games, debates, and negotiations participants may have not fully aligned preferences and may not adopt shared intentions or goals. This paper focuses on modelling negotiations in a multi-issue bargaining setting.

Human-computer negotiation dialogue is typically modelled as a sequence of offers. The offers represent participants' commitments to a certain negotiation outcome. Valuable work has been done on well-structured negotiations - interactions among a few parties with fixed interests and alternatives, see e.g. (Traum et al., 2008a), (Georgila and Traum, 2011), (Guhe and Lascarides, 2014), (Efstathiou and Lemon, 2015). In human negotiation, however, offers as binding commitments are rare and a larger variety of negotiation behavioural patterns is observed (Raiffa et al., 2002a). Participant actions are focused mainly on obtaining and providing preference information and can make this explicitly but also implicitly, see e.g. (Cadilhac et al., 2013). A negotiator often states his preferences without expressing (strong) commitments to accept an offer that includes a positively evaluated option, or to reject an offer that includes a negatively evaluated option.

To achieve more human-like system behaviour, we designed a model which accepts a large variety of dialogue acts representing different levels of commitment. We defined the semantic content of task-related dialogue acts in terms of *negotiation moves*. To model negotiation behav-

ior with respect to preferences, abilities, necessity and acquiescence, and to compute negotiation strategies as accurate as possible, we define several *modal relations* between the modality ‘holder’ (typically the speaker of the utterance) and the target which consists of the negotiation move (and its arguments). Additionally, to facilitate structuring the interaction and enable participants to interpret partner intentions, dynamically changing goals and strategies efficiently, we defined a set of *qualifiers* attached to offer acceptances or rejections and agreements, e.g. tentative or final.

The paper is structured as follows. Section 2 discusses the multi-issue bargaining setting specifying participant tasks, negotiation structure and procedures, actions and other task-related and interactive phenomena observed, and negotiation strategies. In Section 3 the data collection scenario is outlined. Section 4 specifies the dialogue act update semantics. A domain-specific negotiation semantics is discussed in Section 5. We present the performed semantic annotations, and provide corpus statistics. We outline an approach to computing the semantics of negotiation actions. In Section 6 we describe the information state update process in multi-issue bargaining dialogue, leading to the creation of mutual beliefs and belief transfer using various negotiation strategies. Section 7 presents and evaluates the implemented negotiation agent. Section 8 summarises our findings and outlines future research.

2 Multi-issue bargaining

In negotiations, two or more parties have an interest in reaching one or several agreements, and their preferences concerning these agreements are not identical (Raiffa et al., 2002a). *Distributive*, *joint problem-solving* and *integrative* are distinguished.¹ Distributive negotiation means that any gain of one party is made at the expense of the other and vice versa; any agreement divides a fixed pie of value between the parties, see e.g. (Walton and McKersie, 1965). The goal of joint problem-solving negotiations is, by contrast, to work together on an equitable and reasonable solution. The relationship and building up trust is important here (Beach and Connolly, 2005).

In many real-life pure distributive and problem-

¹A fourth type of negotiations is *bad faith*, where parties only pretend to negotiate, but actually have no intention to compromise, see (Cox, 1958)

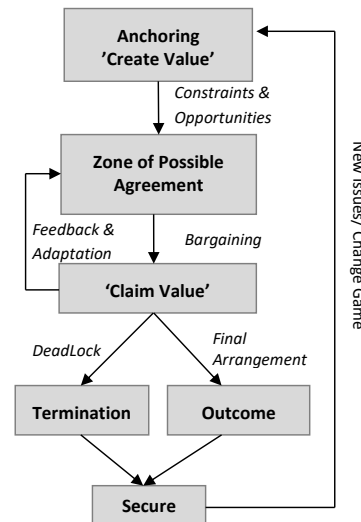


Figure 1: Negotiation phases associated with negotiation structure, based on (Watkins, 2003; Sebenius, 2007).

solving negotiations are rare, more often ‘mixed-motive’ negotiations take place (Lax and Sebenius, 1992). For instance, in sociopolitical and socio-economic contexts, parties are often interested to maintain good long-term relations with each other and therefore try to make trade-offs for both sides to be satisfied with the outcome. At the same time, however, they make competitive efforts to get a bigger share. This problem is often referred to as the ‘Negotiator’s Dilemma’ (Lax and Sebenius, 1992). Negotiators may have partially competitive (adversarial, defensive), and partially cooperative goals. This often happens in integrative multi-issue bargaining, where parties usually have the possibility to simultaneously bargain over several goods and attributes, and to search for integrative potential (interest-based bargaining or win-win bargaining, see Fisher and Ury, 1981).

The different types of negotiation manifest mainly in how parties *create* and *claim values*. Negotiation starts with the **Anchoring** phase, in which participants introduce negotiation issues and options. They also obtain and provide information about preferences, establishing jointly possible values contributing to the **Zone of Possible Agreement** (ZOPA, following the terminology of (Sebenius, 2007)). Participants may bring up early (tentative) offers, typically in the form of suggestions, including referring to the least desirable events - ‘Create Value’. The actual bargaining occurs in the ‘**Claim Value**’ phase, potentially leading to either adaptation, adjustment or can-

celling the originally established ZOPA actions. Patterns of concessions, threats, warnings, and early tentative commitments are observed here. Distributive negotiations are more ‘claiming values’, while joint problem-solving negotiations are more ‘value creating’ interactions, and integrative negotiations are a mix of ‘creating and claiming values’ negotiations (Watkins, 2003). In distributive negotiations the existence and size of the ZOPA is mostly determined by the ‘bottom lines’ of the opposite parties, which are formed by their respective *best alternatives to a negotiated agreement* (BATNA), see (Fisher and Ury, 1981). In integrative bargaining the ZOPA is mainly determined by the number of possible Pareto optimal (or efficient) outcomes. Pareto efficiency (or optimality) reflects a state of affairs when there is no alternative state that would make any partner better off without making anyone worse off.

After establishing the ZOPA, negotiators may still cancel previously made agreements and negotiations might be terminated. **Negotiation Outcome** is the phase associated with the “walk-away” positions for each partner. Finally, negotiators can move to the **Secure** phase summing up, restating reached negotiation agreements or termination outcomes. At this stage, strong commitments are expressed, and weak (mutual) beliefs concerning previously made commitments and reached agreements are strengthened. Participants take decisions to move with another issue or re-start the discussion. Figure 1 depicts the general negotiation structure as described in (Watkins, 2003; Sebenius, 2007) and observed in our data described in the next section.

The outcome of a negotiation depends on the agenda each partner has (Tinsley et al., 2002). The most common tactic of novice negotiators observed is issue-by-issue bargaining (see data collection below). Sometimes, however, negotiators bring all their preferences on the table from the very beginning. This increases the chance to reach a Pareto efficient outcome, since a participant can explore the negotiation space more effectively and reason about each other’s goals, see e.g. (Stevens et al., 2016b). Defensive behaviour, i.e. not revealing preferences, but also being misleading or deceptive, i.e. not revealing true preferences, results in missed opportunities for value creation (Watkins, 2003; Lax and Sebenius, 1992).

The aspects mentioned above may influence ne-

gotiators’ strategies. (Traum et al., 2008a), who studied multi-issue bargaining problem-solving interactions, define strategies as objectives rather than the orientations that lead to them and distinguish seven different strategies: find issue, avoid, attack, negotiate, advocate, success and failure. Other researchers define negotiation strategies closely related to the overall approach for conducting the negotiation and distinguish five strategies: competing (adversarial), collaborating, compromising, avoiding (passive aggressive), and accommodating (submissive), see (Raiffa et al., 2002b; Tinsley et al., 2002). For integrative negotiations, where the negotiators strive for a delicate balance between cooperation and competition, (Lax and Sebenius, 1992), we defined two basic negotiation strategies: cooperative and non-cooperative.

Cooperative negotiators share information about their preferences with their opponents, are engaged in problem-solving behaviours and attempt to find mutually beneficial agreements, (De Dreu et al., 2000). A cooperative negotiator prefers the options that have the highest collective value. If not enough information is available to make this determination, a cooperative negotiator will elicit this information from his opponent. A cooperative negotiator will not engage in positional bargaining² tactics, instead, he will attempt to find issues where a trade-off is possible.

Non-cooperative (or *competitive*) negotiators prefer to withhold their preferences for fear of weakening their power in the negotiation by sharing too much, or they may not reveal true preferences deceiving and misleading the partner. These negotiators focus on asserting their own preferred positions rather than exploring the space of possible agreements, (Fisher and Ury, 1981). A non-cooperative strategy is characterized by a focus on positional bargaining. A negotiator agent using this strategy will rarely ask an opponent for preferences, and will often ignore a partner’s interests and requests for information. A non-cooperative negotiator, instead, will find his own ideal offer, state it, and insist upon it in the hope of making the opponent concede. He will threaten to end the negotiation or will make very small concessions. If the opponent makes a threat, the non-cooperative strategy will accept an offer only if the negotiator can gain a significant number of points from it.

²Positional bargaining involves holding on to a fixed preferences set regardless of the interests of others.

SCOPE	TAXATION
<input type="radio"/> All outdoor smoking allowed	<input type="radio"/> No change in tobacco taxes
<input type="radio"/> No smoking in public transportation	<input type="radio"/> 5% increase in tobacco taxes
<input type="radio"/> No smoking in public transportation and parks	<input type="radio"/> 10% increase in tobacco taxes
<input type="radio"/> No smoking in public transportation, parks and open air events	<input type="radio"/> 15% increase in tobacco taxes
	<input type="radio"/> 25% increase in tobacco taxes
CAMPAIGN	ENFORCEMENT
<input type="radio"/> Flyer and billboard campaign in shopping district	<input type="radio"/> Police fines for minors in possession of tobacco products
<input type="radio"/> Anti-smoking posters at all tobacco sales points	<input type="radio"/> Ban on tobacco vending machines
<input type="radio"/> Anti-smoking television advertisements	<input type="radio"/> Police fines for selling tobacco products to minors
<input type="radio"/> Anti-smoking advertisements across all traditional mass media	<input type="radio"/> Identification required for all tobacco purchases
	<input type="radio"/> Government issued tobacco card for tobacco purchases

Figure 2: Preference card: example of values in four negotiated issues presented in colours.

All the discussed suggests that for adequate modelling we need to take into account that negotiators may perform several types of actions: (1) dialogue acts expressing various levels of commitment; (2) negotiation moves specifying events; (3) qualified ('modalised') actions expressing participants' negotiation strategies; (4) communicative actions to control the interaction.

3 Data Collection

For adequate modelling of human dialogue interactions a systematic analysis of a variety of dialogue phenomena is required. A common procedure of such analysis is human-human data collection and its semantic annotation. The specific setting considered in this study involved a real-life multi-issue bargaining scenario about anti-smoking legislation in the city of Athens passed in 2015-2016. After the new law is enacted, many cases of civil disobedience were reported. Different stakeholders came together to (re-)negotiate and improve the legislation. The main negotiation partner was the Department of Public Affairs of the City Council who negotiates with representatives of small businesses, police, insurances, etc.

The anti-smoking regulations were concerned with four main *issues*: (1) smoke-free public areas (smoking ban scope); (2) tobacco taxes (taxation); (3) effective anti-smoking campaign (campaign); and (4) enforcement policy and police involvement (enforcement), see Figure 2. Each of these issues involves four to five most important negotiation *values* with preferences assigned representing parties negotiation positions, i.e. preference profiles. Nine cases with different preference profiles were designed. The preferences strength was communicated to the negotiators through colours. Brighter orange colours indi-

cated increasingly negative options; brighter blue colours increasingly positive options. The use of colour rather than numbers introduces a form of uncertainty in the exact value of a given agreement, which is closer to real-life negotiations.

Each participant in the experiment received the background story and instructions, as well as their preference profiles. Their task was to negotiate an agreement which assigns exactly one value to each issue, exchanging and eliciting offers concerning $\langle \text{ISSUE}; \text{VALUE} \rangle$ option. Participants were randomly assigned their roles. They were advised to start with the highest possible values according to their preference information. Participants were not allowed to show their preference cards to each other. No further rules on the negotiation process, order of discussed issues or time constraints were imposed. They were allowed to withdraw or re-negotiate previously made agreements within one session, or terminate a negotiation.

16 unique subjects (aged between 19 and 25) participated in the experiments. The resulting data collection consists of 50 dialogues of a total duration of about 8 hours, comprising about 4.000 speaking turns (Petukhova et al., 2016). The human-human negotiation behaviour was evaluated with respect to the number of agreements reached, the ability to find Pareto optimal outcomes, and acceptance of negative outcomes, see Table 6 for results and comparison with human-agent performance. The data was segmented and annotated with dialogue acts information following the ISO 24617-2 standard (ISO, 2012).

4 Dialogue Acts and Update Semantics

In order to model all relevant phenomena, we specified a set of dialogue acts stipulating different levels of commitment with respect to the targeted

Dialogue Act		Relative frequency (in %)	Dialogue Act		Relative frequency (in %)
Communicative function	Modality/ Qualifier		Communicative function	Modality/ Qualifier	
propositionalQuestion		2.0	suggest		10.0
checkQuestion		2.2	addressSuggest		1.4
setQuestion		10.3	acceptSuggest		2.0
choiceQuestion		0.6	declineSuggest		1.7
inform - >		30.3	offer - >		16.7
...	non-modalised	41.3	...	conditional	28.3
...	prefer	30.4	...	tentative	35.0
...	disprefer	3.1	...	final	36.7
...	acquiesce	3.0	addressOffer		0.6
...	need	2.0	acceptOffer - >		5.8
...	able	19.0	...	tentative	47.6
...	unable	1.2	...	final	52.4
agreement		10.3	declineOffer	tentative	2.0
disagreement		4.1			

Table 1: Distribution of task-related dialogue acts in the analysed multi-issue bargaining dialogues.

Dialogue Act	Relative frequency (in %)	
	Task Management	Discourse Structuring
propositionalQuestion	1.8	-
checkQuestion	1.8	-
choiceQuestion	1.8	-
setQuestion	3.5	-
inform	22.8	1.9
answer	7.0	-
(dis-)agreement	10.5	1.9
suggest	22.8	16.8
request	7.0	-
addressSuggest	-	0.9
acceptSuggest	15.8	7.5
declineSuggest	1.8	0.9
offer	1.8	-
addressOffer	1.8	-
interactionStructuring	na	46.7
closing	na	2.8
opening	na	3.7

Table 2: Distribution of Task Management and Discourse Structuring dialogue acts in the analysed multi-issue bargaining dialogues.

negotiated outcome. For this purpose the ISO 24617-2 dialogue act taxonomy³ and its superset DIT⁺⁺⁴, were used. We distinguished five levels of commitment: (1) zero commitment for offer elicitation and preference information requests; (2) the lowest non-zero level of commitment for informing about preferences, abilities and necessities; (3) an interest and consideration to offer a certain value; (4) weak (tentative) or conditional commitment to offer a certain value; and (5) strong (final) commitment to offer a certain value.

Actions at zero level of commitment are used by negotiators to gather information about partner’s preferences, mostly in the form of questions. For example, a *Set Question* of participant *A* addressed to *B* with the goal to elicit *B*’s preference concerning the smoking ban scope, e.g. ‘Where do you think we should ban smoking?’, can

be represented as $SetQuestion(A, B, offer(ISSUE = 1; ?VALUE))$. To describe the intended update effects of an action a number of formal concepts - semantic primitives - are used that specify an agent’s beliefs, goals, and commitments. A set of semantic primitives is defined in (Petukhova, 2011). Bunt (2014) provides a detailed specification of the update semantics of dialogue acts. For instance, the primitive *Bel* expresses the possession of information and the *KnowVal* primitive serves to represent the availability of information. For example, *A* believing that *B* has certain preferences for the ‘scope’ issue is represented as $Bel(A, KnowVal(B, offer(ISSUE = 1; ?VALUE)))$.⁵ The primitive *Want* is used to capture a participant’s goal to achieve a certain situation. Thus, *A*’s goal to obtain information about *B*’s negotiation preference can be represented as $Want(A, KnowVal(A, offer(ISSUE = 1; ?VALUE)))$.

Negotiators may *Inform* each other about their preferences. These actions also include various types of *Answers*. *A*’s goal to inform *B* about his negotiation preferences can be represented as $Want(A, KnowVal(B, Bel(A, offer(ISSUE = 1; VALUE = 1C)))$. Negotiators do not just provide information about their preferences, but also communicate their evaluation and estimation of the probability of events and their beliefs about what is possible, necessary and desirable in the current context, e.g. $Bel(A, \square offer(ISSUE = 1; VALUE = 1C))$, $Bel(A, \neg \diamond offer(ISSUE = 1; VALUE = 1C))$.

Suggestion acts express considerations to offer certain values, and assumptions about the oppo-

³See http://dit.uvt.nl/#iso_24617-2

⁴<http://dit.uvt.nl/>

⁵Additionally, the strength of *A*’s beliefs is represented by the parameter σ , which can have the values ‘firm’ and ‘weak’, or numerical values expressed, for example, by confidence scores computed elsewhere.

ment’s abilities and interests to offer the same, i.e. $ConsidDo(A, offer(X;Y)); Bel(A, CanDo(B, offer(X;Y)))$; $Bel(A, Interest(B, offer(X;Y)))$.

At a higher level of commitment, *Offer* acts are observed, expressing commitments to offer (or not to offer) a certain value, e.g. $CommitDo(A, offer(X;Y))$ and $CommitRefrain(A, offer(X;Y))$. Weak and strong commitments to *Accept* or *Reject* an *Offer* (but also a *Suggestion*) may dependent on a condition specified in the semantic content of a dialogue act. When annotating, a *conditional* qualifier is attached to action-discussion communicative functions. Additionally, all offers and responses to them at all negotiation stages except the *Secure* phase are modelled as weak commitments, e.g. $WBel(A, CommitDo(A, offer(X;Y)))$, indicating that they are *tentative*, and can eventually be strengthened or cancelled. At the highest level of commitment are *final offers* and responses to them in the *Secure* phase. Annotations contain tentative and final communicative function qualifiers. Table 1 presents the observed distribution of task-related dialogue acts in the annotated data.

To structure a negotiation task, *Task Management* acts are used. These dialogue acts explicitly address the negotiation process and procedure. This includes utterances for coordinating the negotiators’ activities (e.g., “Let’s go issue by issue”) or asking about the status of the process (e.g., “Are we done with the agenda?”). Task Management acts are specific for a particular task and are often similar in form but different in meaning from Discourse Structuring acts, which address the management and monitoring of the interaction. Examples of the later are utterances like “To sum up”, and “Let’s move to a next round”. Table 2 presents the distribution of these dialogue acts.

5 Negotiation Semantics

Semantically, dialogue acts correspond to update operations on the information states of the dialogue participants. They have two main components: (1) the *communicative function*, that specifies how to update an information state, e.g. *Inform*, *Question*, and *Request*, and (2) the *semantic content*, i.e. the objects, events, situations, relations, properties, etc. involved in the update, see (Bunt, 2000). Negotiations are commonly analysed in terms of certain actions, such as offers, counter-offers, and concessions, see (Watkins,

Negotiation Move	Relative frequency (in %)
Offer	75.0
CounterOffer	12.4
Exchange	6.6
Concession	1.2
BargainIn	0.4
BargainDown	1.2
Deal	2.4
Withdraw	0.8

Table 3: Defined negotiation moves and their relative frequencies in the annotated multi-issue bargaining corpus.

2003), (Hindriks et al., 2007). We considered two possible ways of using such actions, also referred to as ‘negotiation moves’, to compute the update semantics in negotiation dialogues. One is to treat negotiation moves as task-specific dialogue acts. Due to its domain-independent character, the ISO 24617-2 standard does not define any communicative functions that are specific for a particular kind of task or domain, but the standard invites the addition of such functions, and includes guidelines for how to do so. For example, a negotiation-specific kind of $Offer_N$ function should be introduced for the expression of commitments concerning a negotiation value.⁶ Another possibility is to use negotiation moves as the semantic content of general-purpose dialogue acts. For example, a negotiator’s statements concerning his preference to a certain option can be represented as $Inform(A, B, \diamond offer(X;Y))$.

We specified 8 basic negotiation moves, see distribution in the analysed data in Table 3.

Negotiators often communicate their cooperativity by using modal utterances expressing preference and ability. Non-cooperative behaviour, by contrast, may be articulated by expressing inability and dislike. Modality expressions are mainly observed in *Inform* and *Answer* acts, see Table 1.

The proposed approach allows for flexibility in the interpretation and generation of negotiation strategies and accounts for a richer set of task-related actions.

6 Belief Transfer and Negotiation Strategies

We compute the meaning of negotiation dialogue contributions in terms of their effects on the participants’ information states as proposed by the Information State Update (ISU) approach, (Poesio and Traum, 1998; Bunt, 1989) and the com-

⁶Negotiation ‘Offers’ may have a more domain-specific name, e.g. *Bid* for selling-buying bargaining.

Context	num	source	Agent (A)context	num	source	Council (C) context
LC				u001	<i>prec</i>	<i>Bel(C,Next_Speaker(C))</i>
LC	s1 <i>fs</i> ₁ <i>da</i> ₁	<i>latest</i> <i>D;CF</i> <i>sem.content</i>	<i>Bel(A,Current_Speaker(C))</i> (<i>verbatim</i>) Task; Suggest <i>p2 = offer(ISSUE = 1;VALUE = 1b)</i> Speaker:C; Addressee: A	u1 <i>fs</i> ₁ <i>da</i> ₁	<i>latest</i> <i>D;CF</i> <i>sem.content</i>	<i>Bel(C,Current_Speaker(C))</i> (<i>verbatim</i>) Task; Suggest <i>p2 = offer(ISSUE = 1;VALUE = 1b)</i> Speaker:C;Addressee: A
CC	s2	exp.und: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C,Interpreted(A,du1))))</i>	u2	exp.und: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C,Interpreted(A,du1))))</i>
SC	s01a	exp.und: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C, Bel(A,Bel(C,Interest(A,p2))))))</i>	u01a	exp.und: <i>da</i> ₁	<i>Bel(C,MBel({A,C},WBel(C, Bel(A,Bel(C,Interest(A,p2))))))</i>
	s01b	exp.und: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C, Bel(A,Assume(C,CanDo(A,p2))))))</i>	u01b	exp.und: <i>da</i> ₁	<i>Bel(C,MBel({A,C},WBel(C, Bel(A,Assume(C,CanDo(A,p2))))))</i>
	s01c	exp.und: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C, Bel(A,Want(C,ConsidDo(A,p2))))))</i>	u01c	exp.und: <i>da</i> ₁	<i>Bel(C,MBel({A,C},WBel(C, Bel(A,Want(C,ConsidDo(A,p2))))))</i>
	s02a	exp.ad: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C, Bel(A,Interest(A,p2))))))</i>	u02a	exp.ad: <i>da</i> ₁	<i>Bel(C,MBel({A,C},WBel(C, Bel(A,Interest(A,p2))))))</i>
	s02b	exp.ad: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C, Bel(A,CanDo(A,p2))))))</i>	u02b	exp.ad: <i>da</i> ₁	<i>Bel(C,MBel({A,C},WBel(C, Bel(A,CanDo(A,p2))))))</i>
	s02c	exp.ad: <i>da</i> ₁	<i>Bel(A,MBel({A,C},WBel(C, Bel(A,ConsidDo(A,p2))))))</i>	u02c	exp.ad: <i>da</i> ₁	<i>Bel(C,MBel({A,C},WBel(C, Bel(A,ConsidDo(A,p2))))))</i>
SC	s03a	und: <i>da</i> ₁	<i>Bel(A,Bel(C,Interest(A,p2)))</i>			
	s03b		<i>Bel(A,Assume(C,CanDo(A,p2)))</i>			
	s03c		<i>Bel(A,Want(C,ConsidDo(A,p2)))</i>			
SC	s3	<i>prec</i>	<i>Bel(A,◇p2)</i>			
	s04a	ad: <i>da</i> ₁	<i>Bel(A,Interest(A,p2))</i>			
	s04b		<i>ConsidDo(A,p2)</i>			
	s04c		<i>Bel(A,CanDo(A,p2))</i>			
SC	s4	<i>prec</i>	<i>CommitDo(A,p2)</i>			
LC	<i>da</i> ₂	<i>plan:s4</i> <i>sem.content</i>	Task; AcceptSuggest <i>p2 = offer(ISSUE = 1;VALUE = 1b)</i>			
LC	s001	<i>prec</i>	<i>Bel(A,Next_Speaker(A))</i>			
LC	s5 <i>fs</i> ₂ <i>da</i> ₂	<i>latest</i> <i>D;CF</i> <i>sem.content</i>	<i>Bel(A,Current_Speaker(A))</i> (<i>verbatim</i>) Task;AcceptSuggest <i>p2 = offer(ISSUE = 1;VALUE = 1b)</i> antecedent: <i>da</i> ₁ Speaker:A; Addressee: C	u2 <i>fs</i> ₂ : <i>du</i> ₂ <i>fs</i> ₂ : <i>da</i> ₂	<i>latest</i> <i>D;CF</i> <i>sem.content</i>	<i>Bel(C,Current_Speaker(A))</i> (<i>verbatim</i>) Task;AcceptSuggest <i>p2 = offer(ISSUE = 1;VALUE = 1b)</i> antecedent: <i>da</i> ₁ Speaker:A;Addressee: C

Table 4: Example of context update for cooperative negotiation behaviour. (LC = Linguistic Context; CC = Cognitive Context; SC = semantic context; *prec* = preconditions; *da* = dialogue act; *fs* = functional segment; *D* = dimension; *CF* = communicative function; exp.und = expected understanding; und = understanding; exp.ad = expected adoption; ad = adoption; *Bel* = believes; *MBel* = mutually believed; *WBel* = weakly believes)

putational model of grounding and beliefs transfer proposed by (Bunt et al., 2007).

Negotiators aim at the understanding by others. Understanding that a certain dialogue act is performed means creating the belief that the preconditions which are characteristic for that dialogue act hold. Using the ISU machinery and procedures for incorporating beliefs and expectations shared between speaker and hearers, we can compute *expected understanding effects* modelled as *weak* beliefs. When evidence about successful understanding arrives weak beliefs are *strengthened*, otherwise they may be *cancelled*.

Negotiators also expect that their opponent will accept some of their offers (*expected adoption effects*). The strength of such expectations depends on their knowledge about their opponents, on their goals, and on the knowledge concerning the opponent's negotiation strategy. When the negotiator states identical preferences, agrees with the opponent's preferences, or accepts his suggestions and offers, he adopts the opponent's beliefs as beliefs of his own. For example:

- (1) Council(human): *What do you think if we do not allow smoking in public transportation at least?*

Business(agent): *Well, I think we can live with that*

Council (C) produces a $\langle Task; suggest \rangle$ dialogue act with the semantic content *p2*. Weak mutual beliefs concerning expected understanding and adoption effects are created, the dialogue context model is updated with *s01a – s02c* and *u01a – u02c* updates as shown in Table 4. Business representative A understands C's *da*₁ as a suggestion and accepts it following the cooperative negotiation strategy. A's understanding means that A believes that C wants A to consider to do *p2* because C believes that *p2* would be interesting for A, and A is able to do *p2*. In A's preference profile, *p2* is a possible offer. This enables A to accept C's suggestion, see precondition in *s3*. A acting as a cooperative agent is considering to offer the discussed value and commits to perform this action. Thus, beliefs about expected and actual understanding and adoption together with the negotiator's preferences give rise to the generation of one or more relevant dialogue acts. Similarly, additional updates are performed in other contexts. For instance, the Linguistic Context (LC) is updated with respect to beliefs concerning the speaker role management, and in the Cognitive Context (CC) concerning processing successes and failures. This triggers the gener-

ation of dialogue acts in multiple dimensions, e.g. here in the Turn Management and Feedback dimensions, respectively.

The example in (2) shows non-cooperative negotiation behaviour. It may be noted that negotiation partners always cooperate at a linguistic level, as they try to understand each other's contributions and respond to perceived intentions.⁷ A rational agent may show non-cooperative behavior at the level of perlocutionary actions (see (Attardo, 1997)), when cancelling of expected adoption beliefs occurs.

- (2) Council(human): *What do you think if we do not allow smoking in public transport at least?*
 Business(agent): *It's not possible for me*

The dialogue context model is updated in this case as follows. A understanding C means that A believes that C wants A to consider to do $p2$ because C believes that $p2$ would be interesting for A and A is able to do $p2$. According to A 's preference profile, $p2$ is not a possible offer, resulting in the precondition in $s3$ as $Bel(A, \neg \diamond p2)$. This leads to cancelling C 's expected adoption beliefs. Acting as a non-cooperative but rational agent, A refuses to commit to $p2$. Alternatively, A may offer another value more preferable for him, i.e. performing a counter-offer move when $Bel(A, Interest(A, \neg p2))$ but $Bel(A, Interest(A, p3)); Bel(A, CanDo(A, offer(p3))); ConsidDo(A, offer(p3))$ where $p3$ stands for example for $offer(ISSUE = 1; VALUE = 1c)$.

7 Negotiation Agent

The implemented Negotiation Agent produces counter-move based on the estimation of partner's preferences and goals. The Agent adjusts its strategy according to the perceived level of the opponent's cooperativeness. Such meta-strategies are observed in human negotiation and coordination games, see (Kelley and Stahelski, 1970), (Smith et al., 1982). Currently, the Agent distinguishes three strategies: cooperative, non-cooperative and neutral. The agent starts neutrally, requesting the partner's preferences. If the Agent believes the opponent is behaving cooperatively, it will react with an cooperative negotiation move. For instance, it will reveal its preferences when asked for, it will accept

⁷Consider also the definition of cooperative communicative behaviour proposed by Allwood et al., 2000. Communicative agents are considered cooperative at least in trying to recognise each other's goals, and the recognition of a goal may be sufficient reason for the participant to form the intention to act.

the opponent's offers, and propose concessions or cross-issues trade-offs. It will use modality triggers of liking and ability. If the Agent experiences the opponent being non-cooperative, it will switch to non-cooperative mode. It will stick to its preferences and insist on acceptance by the opponent. It will repeatedly reject the opponent's offers using modal expressions of inability, dislike and necessity. It will not make concessions and will threaten to withdraw reached agreements and/or terminate negotiation.

The Agent's negotiation moves and their arguments are encoded as instances represented as a set of slot-value pairs corresponding to the Agent's preference profile concerning beliefs about Agent's and partner's preferences (state of the negotiation and conditions), and Agent's and estimated partner's goals (actions), see Table 5. The Agent assumes that the partner's preferences are comparable, but values may differ. At the beginning of the interaction, the Agent may have no or weak assumptions (guesses) about the partner's preferences. As the interaction proceeds the Agent builds up (learns) more knowledge about his partner's choices. Agent's decisions are made by finding a prior experience (an instance) that is most 'active' (based on history, e.g. frequency and recency, and on similarity, e.g. how similar the instance is, given the context) in the current context, see (Gonzalez and Lebiere, 2005). Negotiation Agent is based on the Instance-Based Learning (IBL) model as implemented in ACT-R cognitive architecture (Anderson, 2007).

Having computed the 'best' negotiation move as a response, the Agent will pass it to the Dialogue Manager for updating the dialogue context model and producing an appropriate task-related dialogue act. Thus, the Negotiation Agent is integrated in a spoken dialogue system as a Task Agent of its Dialogue Manager which operates on a structured dynamic dialogue context, see (Malchanau et al., 2015) for the proposed multi-threaded DM architecture.

We evaluated the Negotiation Agent's performance, comparing it with human performance on the number of agreements reached, the ability to find Pareto optimal outcomes, the degree of cooperativeness, and negative outcomes, see Table 6. For this evaluation, 28 sessions involving 28 participants aged 25-45 (all professional politicians or governmental workers) were analysed.

Information type	Explanation	Source
Strategy	The strategy associated with the instance	negotiationMove, modality
My-bid-value-me	The number of points the agent’s bid is worth to the agent	
My-bid-value-opp	The number of points that the agent believes its bid is worth to the user	
Opp-bid-value-me	The number of points the users bid is worth to the agent	
Opp-bid-greater	<i>true</i> if the users bid is at least as much as the agent’s current bid, <i>false</i> otherwise	Preference profile
Next-bid-value-me	The number of points that the next best option is worth The next best option is defined as the option closest in value to the current one (Not including those that are worth more than the current option.)	
Overall-value	The total value of all options that have been agreed upon so far. This is a measure of how the negotiation is going. If it is negative, negotiation is likely to result in an unacceptable outcome.	History
My-move	The move that the agent should take in this context.	Planned future

Table 5: Structure of an instance in the Negotiation Agent, adopted with extensions from (Stevens et al., 2016a).

Evaluation criteria	Human-human	Human-computer
Mean dialogue duration (min)	5:51	9:37
Agreements (%)	78	66
Pareto optimal (%)	61	60
Negative deal (%)	21	16
Cooperativeness rate (%)	39	51

Table 6: Comparison of human-human and human-agent negotiation behaviour.

We found that the participants reached a lower number of agreements when negotiating with the Agent than when negotiating with each other. This could in most cases be attributed to the imperfect recognition and interpretation by the dialogue system of spoken participant behaviour. Overall task effectiveness in terms of proportion of successfully completed dialogues was found to be 76.8% (human-human pairs were 100% successful). Of the reached agreements, the participants made a similar number of Pareto optimal agreements when negotiating with the Agent as when negotiating with each other. Human participants show a higher level of cooperativity when interacting with the Agent, measured in the number of cooperative actions given the total number of the task-related actions performed. This may mean that humans were more competitive when interacting with each other. A lower number of negative deals (i.e. agreements on bright ‘orange’ options in Figure 2) was observed for human-agent pairs.

8 Conclusions and Future Research

In this study we proposed, implemented and evaluated an ISU-based model of multi-issue bargaining dialogue behaviour. A real-life complex negotiation scenario was used for data collection, with a rather comprehensive pragmatic and semantic analysis of negotiation phenomena. The model accounts for specific multi-issue bargaining dialogue

structure, for actions that express different degrees of commitment to targeted negotiation outcome, as well as for the strategic actions to achieve this outcome. The model is flexible in that it can be extended with other domain-specific event-based semantics. We showed how the participants’ beliefs are created when a speaker’s behaviour is understood and how it leads to the adoption or cancellation of beliefs when participants have overlapping and conflicting preferences. The model supports the generation of dialogue contributions in multiple dimensions accounting for task-related negotiation actions as well as for actions that are used to control the overall interaction.

The evaluation of human-human and human-agent performance shows that the relevant negotiation aspects and interactive phenomena are adequately modelled, resulting in plausible and effective negotiation behaviour.

Future efforts will be undertaken to refine the model with respect to the negotiation moves semantics. We also plan to extend the model to account for attitudinal meaning aspects of multimodal dialogue contributions to compute sophisticated negotiation strategies with respect to cooperativity and dominance. A user-based within-subject evaluation (e.g. in repetitive negotiation rounds) will be performed to analyse participant’s negotiation behaviour change over time, and to incorporate user models into the adaptive human-computer negotiation system.

Acknowledgments

This research is partly funded by the EU FP7 Metalogue project, under grant agreement number: 611073. We are also very thankful to anonymous reviewers for their valuable comments.

References

- Jens Allwood, David Traum, and Kristiina Jokinen. 2000. Cooperation, dialogue and ethics. *International Journal of Human Computer Studies*, pages 871–914.
- John R. Anderson. 2007. *How Can the Human Mind Occur in the Physical Universe?* New York, NY: Oxford University Press.
- Nicholas Asher and Jason Quinley. 2011. Begging questions, their answers and basic cooperativity. In *JSAI International Symposium on Artificial Intelligence*, pages 3–12. Springer.
- Salvatore Attardo. 1997. Locutionary and perlocutionary cooperation: The perlocutionary cooperative principle. *Journal of Pragmatics*, 27(6):753–779.
- Lee R Beach and Terry Connolly. 2005. *The psychology of decision making: People in organizations*. Sage.
- Harry Bunt, Simon Keizer, and Roser Morante. 2007. A computational model of grounding in dialogue. In *Proceedings of the Workshop in Discourse and Dialogue. Lecture Notes in Computer Science 4629*, pages 591–598, Antwerp, Belgium.
- Harry Bunt. 1989. Information dialogues as communicative action in relation to partner modelling and information processing. In M. Taylor, F. Neel, and D. Bouwhuis, editors, *The Structure of Multimodal Dialogue*, volume 1, pages 47–73. Elsevier, North Holland, The Netherlands.
- Harry Bunt. 2000. Dialogue pragmatics and context specification. In H. Bunt and W. Black, editors, *Abduction, Belief and Context in Dialogue; studies in computational pragmatics*, pages 81–105. John Benjamins, Amsterdam.
- Harry Bunt. 2014. A context-change semantics for dialogue acts. In *Computing meaning*, pages 177–201. Springer.
- Anais Cadilhac, Nicholas Asher, Farah Benamara, and Alex Lascarides. 2013. Grounding strategic conversation: Using negotiation dialogues to predict trades in a win-lose game. In *EMNLP*, pages 357–368.
- Archibald Cox. 1958. The duty to bargain in good faith. *Harvard Law Review*, pages 1401–1442.
- Carsten KW De Dreu, Laurie R Weingart, and Seungwoo Kwon. 2000. Influence of social motives on integrative negotiation: a meta-analytic review and test of two theories.
- Ioannis Efstathiou and Oliver Lemon. 2015. Learning non-cooperative dialogue policies to beat opponent models: the good, the bad and the ugly. *SEMDIAL 2015 goDIAL*, page 33.
- Roger Fisher and William L Ury. 1981. *Getting to yes: Negotiating agreement without giving in*. Harmondsworth, Middlesex: Penguin.
- Kallirroi Georgila and David Traum. 2011. Reinforcement learning of argumentation dialogue policies in negotiation. In *Twelfth Annual Conference of the International Speech Communication Association*.
- Cleotilde Gonzalez and Christian Lebiere. 2005. Instance-based cognitive models of decision-making. In D. Zizzo and A. Courakis, editors, *Transfer of knowledge in economic decision making*. Macmillan.
- Markus Guhe and Alex Lascarides. 2014. Persuasion in complex games. *DialWattSemdial 2014*, page 62.
- Koen Hindriks, Catholijn M Jonker, and Dmytro Tykhonov. 2007. Analysis of negotiation dynamics. In *International Workshop on Cooperative Information Agents*, pages 27–35. Springer.
- ISO. 2012. *Language resource management – Semantic annotation framework – Part 2: Dialogue acts. ISO 24617-2*. ISO Central Secretariat, Geneva.
- Harold H Kelley and Anthony J Stahelski. 1970. Social interaction basis of cooperators’ and competitors’ beliefs about others. *Journal of personality and social psychology*, 16(1):66.
- David Lax and James Sebenius. 1992. The manager as negotiator: The negotiators dilemma: Creating and claiming value. *Dispute resolution*, 2:49–62.
- Andrei Malchanau, Volha Petukhova, Harry Bunt, and Dietrich Klakow. 2015. Multidimensional dialogue management for tutoring systems. In *Proceedings of the 7th Language and Technology Conference (LTC 2015)*, Poznan, Poland.
- Volha Petukhova, Christopher Stevens, Harmen de Weerd, Niels Taatgen, Fokie Cnossen, and Andrei Malchanau. 2016. Modelling multi-issue bargaining dialogues: data collection, annotation design and corpus. In *Proceedings 9th International Conference on Language Resources and Evaluation (LREC 2016)*. ELRA, Paris.
- Volha Petukhova. 2011. *Multidimensional Dialogue Modelling. PhD Thesis*. Tilburg University, The Netherlands.
- Massimo Poesio and David Traum. 1998. Towards an axiomatization of dialogue acts. In *Proceedings of the Twente Workshop on the Formal Semantics and Pragmatics of Dialogues*, pages 207–222.
- Howard Raiffa, John Richardson, and David Metcalfe. 2002a. *Negotiation analysis: The science and art of collaborative decision making*. Harvard University Press.

- Howard Raiffa, John Richardson, and David Metcalfe. 2002b. *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Cambridge: Belknap Press.
- James K Sebenius. 2007. 23 negotiation analysis: Between decisions and games. *Advances in Decision Analysis: From Foundations to Applications*, page 469.
- D Leasel Smith, Dean G Pruitt, and Peter J Carnevale. 1982. Matching and mismatching: The effect of own limit, other's toughness, and time pressure on concession rate in negotiation. *Journal of Personality and Social Psychology*, 42(5):876.
- Christopher A Stevens, Harmen de Weerd, Fokie Cnossen, and Niels A Taatgen. 2016a. A metacognitive agent for training negotiation skills. In *Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016)*.
- Christopher A Stevens, Niels A Taatgen, and Fokie Cnossen. 2016b. Instance-based models of metacognition in the prisoner's dilemma. *Topics in cognitive science*, 8(1):322–334.
- Catherine H Tinsley, Kathleen M O'Connor, and Brandon A Sullivan. 2002. Tough guys finish last: The perils of a distributive reputation. *Organizational Behavior and Human Decision Processes*, 88(2):621–642.
- David Traum, Stacy C Marsella, Jonathan Gratch, Jina Lee, and Arno Hartholt. 2008a. Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In *International Workshop on Intelligent Virtual Agents*, pages 117–130. Springer.
- David Traum, William Swartout, Jonathan Gratch, and Stacy Marsella. 2008b. A virtual human dialogue model for non-team interaction. In *Recent trends in discourse and dialogue*, pages 45–67. Springer.
- Richard E Walton and Robert B McKersie. 1965. *A behavioral theory of labor negotiations: An analysis of a social interaction system*. Cornell University Press.
- Michael Watkins. 2003. Analysing complex negotiations. *Harvard Business Review*.