

Negotiating Agents

*Catholijn M. Jonker, Koen V. Hindriks,
Pascal Wiggers, Joost Broekens*

■ *Negotiation is a complex emotional decision-making process aiming to reach an agreement to exchange goods or services. From an agent technological perspective creating negotiating agents that can support humans with their negotiations is an interesting challenge. After more than a decade of research, negotiating agents can outperform human beings (in terms of deal optimality) if the negotiation space is well understood. However, the inherent semantic problem and the emotional issues involved mean that negotiation cannot be handled by artificial intelligence alone, and a human-machine collaborative system is required. This article presents research goals, challenges, and an approach to create the next generation of negotiation support agents.*

Everyone is an experienced negotiator and everyone has an opinion about his or her negotiation skills. However, even professional negotiators can still improve their skills considerably. “Most people are ineffective negotiators Fewer than 4 percent of managers reach win-win outcomes when put to the test Even on issues for which people were in perfect agreement, they fail to realize it 50 percent of the time” (Thompson 2005). Although many forms of negotiation exist, in this article we focus on integrative bargaining (see Walton and McKersie 1965).

Negotiation is a prime example of a task for which the human mind is but partially equipped, and for which artificial intelligence (AI) can provide assistance. Among others AI can provide search techniques, computational heuristics to tackle exponential problem spaces, strategic reasoning, argumentation, learning techniques, and affective computing to handle the complications that arise in negotiations. More difficult problems that are not as easily solved by artificial intelligence techniques, however, include obtaining the common knowledge necessary to understand negotiation domains and arbitrary human conversations that take place during negotiations.

We aim for synergy between human and agent in such a way that the human weaknesses are covered by the strengths of the agent and the weaknesses of the agent are covered by the strengths of the human. This implies that tasks should be divided over humans and agents in a way that respects those capabilities. On the one hand, humans are better equipped to understand the context and the emotional fluctuations in human-human interaction, they are capable of finding new

relations between concepts, and they have the necessary background knowledge to interpret the domain of negotiation with respect to their preferences. On the other hand, humans can be troubled by emotions, and have difficulty handling the complexity of negotiation spaces and keeping track of the interests of the negotiation opponent. For agents it is largely the other way around.

The rest of this article is organized as follows. We first informally introduce human negotiation as a process of four phases that may be distinguished. We then proceed by discussing some state-of-the-art negotiation support systems and automated negotiating agents. We then present the pocket negotiator project that is developing the first of a next generation of negotiation support systems. Thereafter we discuss various technical components based on different artificial intelligence techniques that are part of this agent, including support for analyzing a negotiation, taking the opponent into account, advice on how to negotiate strategically, for eliciting human preferences, and for handling emotions. The article concludes with an overview of open research questions.

Negotiation in Phases

Fisher, Ury, and Patton (1992); Raiffa, Richardson, and Metcalfe (2002); Thompson (2005); and others emphasize that negotiation is not just about money, but also about good relationships, awareness of all issues being negotiated, personal preferences of both parties, knowledge of your alternatives (if no deal is reached), and reflection on your performance. Negotiation is a process that is subject to cultural differences; see Hofstede, Jonker, and Verwaart (2010). Although the number of stages in the literature varies, the following four major stages can be discerned in integrative negotiation: private preparation, joint exploration, bidding, and closing.

Private preparation is a stage of information gathering and reflection done before meeting the other party. The negotiator learns as much as possible about the negotiation domain (issues under negotiation and hidden interests), the coming process, about his/her profile, and about the opponent. Hidden interests are aspects that might not be mentioned but that do have an impact; for example, is one of the parties under time pressure? During this phase a machine can most effectively assist a human with exploring his/her preferences and getting a realistic picture of the negotiation possibilities; for example, what is the current market price?

In the *joint exploration* phase the negotiating parties talk to each other but don't place bids on the table. The aim of this stage is to check the information they gathered so far, to create a good atmosphere for the bidding that will follow, to

make the negotiation space as big as possible, and to agree upon a protocol for the bidding, for example, turn-taking by phone. In the joint exploration phase, the machine can assist a human in refining the domain of negotiation and constructing a model of his or her opponent.

During the bidding stage negotiators exchange bids. There are two key strategic considerations a negotiator has to make in this phase. First, a negotiator needs a bidding strategy to determine the next bid. Second, a negotiator needs an acceptance strategy to decide whether to accept an incoming bid (if he or she expects no more improvements can be made), to make a counteroffer (if the negotiator thinks he or she can do better), or to stop (if the negotiator thinks he or she has a better alternative elsewhere). In the bidding phase, the machine can support a human by advising or teaching it which bid to make.

During the *closing stage* the outcome of the bidding stage is formalized and confirmed by both parties. If confirmation turns out to be impossible, the negotiation returns to one of the previous stages. Sometimes it is beneficial to enter a postnegotiation phase to explore whether the agreement reached can be improved upon by revealing the preferences of both negotiating parties. In this phase, machines can help determine whether an agreement can be improved upon.

Overall for humans, negotiating is an emotional process; certainly for the novice negotiator (Ury 1993, 2007). The more that depends on the outcome of the negotiation, the more intense the emotions. For example, buying a house for the first time, or negotiating about a job contract, can be intense; see for example, DeLuca and DeLuca (2007). This is partly caused by not feeling in control of the situation, not knowing what to expect, and fearing not to perform well enough (Folkman and Lazarus 1990, Lazarus and Folkman 1984, Ursin and Erisen 2004).

The human perspective on negotiation is central in this article. The next sections discuss some state-of-the-art negotiation support systems (NSS) and some challenges of creating negotiation support. However, the article also briefly discusses the solid repository of automated negotiating agents (ANA) and their bidding and acceptance strategies for integrative bargaining.

Negotiation Support: State of the Art

There are only a few negotiation support systems that are used in practice: Inspire, Athena, and SmartSettle.

Inspire¹ is a web-based negotiation support system. It contains a facility for specification of preferences and assessment of offers, an internal mes-

saging system, graphical displays of the negotiation's progress, and other capabilities. It has been used to support humans in negotiation as well as to collect data about such negotiations for research purposes. It offers the user a structured approach to prepare and engage in a negotiation, and can be used as a training tool.

Another NSS example is provided by Athena², which has been primarily used in education. As is the case for Inspire, users of Athena have to build the domain models themselves. That is, preferences are elicited from the user who has to provide the domain structure. The support does not include predefined repositories of domain models, interaction support, or assistance in selecting a bidding strategy.

Smartsettle³ is a commercial negotiation support system that also provides bidding support. Interestingly, while other systems keep offers and demands hidden, Smartsettle displays proposals and suggestions to all parties. It is strong in mediation, and not developed for closed integrative bargaining.

The systems here described are high-quality systems that have proved their worth in practice and, furthermore, showed a need for negotiation support. These systems help users to structure negotiations and help researchers to understand human difficulty in handling negotiations. The systems do not provide training and don't prepare the user for interaction with an opponent or for understanding the role of emotions in a negotiation. They also do not provide advanced bidding support that takes the negotiation style of the user into account, nor do they provide advice on when to accept a bid in closed integrative bargaining. Finally, these systems lack the intelligence to learn and estimate the preference profile and strategies of the opponent and to learn average profiles from all negotiators that negotiate about the same preference profile. The Pocket Negotiator project attempts to create a negotiation support system that addresses most of these shortcomings.

Automated Negotiating Agents

Over the past decade various models for automated negotiating agents have been proposed and many results on the performance of such agents have been published (Jonker and Treur 2001; Meyer et al. 2004; Rahwan, Sonenberg, and McBurney 2005; Büttner 2006; Hindriks, Jonker, and Tykhonov 2007). The research has mainly focused on devising strategies, protocols, and negotiation languages, that is, languages to represent negotiation domains (Rosenschein and Zlotkin 1994, Kraus 2001, Tamma et al. 2005). Among others, it has been demonstrated and replicated that automated negotiating agents may obtain significant

improvements over the outcomes obtained by humans (see, for example, Bosse and Jonker [2005]). Additionally, learning techniques have been developed to learn the preferences or the strategy of the other party (see, for example, Oliver 2005). Such techniques may also be useful for eliciting preferences. The state of the art in automated negotiating agents is presented at the newly formed yearly Automated Negotiating Agents Competition (ANAC) (Ito et al. 2010).⁴

The primary purpose of the negotiation competition is to stimulate research about automated agents for bilateral, multiissue negotiation. Purely analytic methods based on, for example, game-theoretic approaches cannot be directly applied to design efficient negotiating agents due to incomplete information about opponents and the generally complex multiissue domains. The competition offers a venue for evaluating the strengths and weaknesses of negotiating agents that use heuristics to deal with these complications. Agents are evaluated based on their ability to obtain high utility outcomes. As it is well known that strategies may perform differently on different negotiation domains, agents play on a range of different domains against each other in a full tournament. An additional benefit of the yearly competition is that it helps build a best practice repository of negotiation techniques.

The automated negotiating agents competition aims for design of more efficient negotiating agents; testing bidding and acceptance strategies; exploring learning strategies and opponent models; and collecting the state-of-the-art negotiating agents, negotiation domains, and preference profiles and making them available for the negotiation research community and related communities.

To facilitate research on bilateral multiissue negotiation and to be able to run negotiation tournaments, the Genius system has been developed.⁵ It allows easy development of negotiating agents and integration of existing negotiating agents.

The Pocket Negotiator

Our aim is to develop a Pocket Negotiator agent that assists (not supplants) the user in an integrative bargaining task: assessing the situation, regulating emotions, and coping with negative consequences of emotions. We divide the tasks between the user and the agent in accordance to the strengths and weaknesses of both. To ensure optimal teamwork, user and agent need to share a generic model of negotiation (Brazier et al. 2000; Jonker, Riemsdijk, and Vermeulen 2011) and of their respective strengths and weaknesses. Human strengths are the wealth of general world knowledge and their communication proficiency. Agents, however, can be equipped with specialised

knowledge about negotiation, emotions, and specific negotiation domains. Furthermore, agents can improve the utility of an agreement.

The Pocket Negotiator agent is to enhance the negotiation skills and performance of the user by helping the user to explore the negotiation space and obtain win-win outcomes that are better for both parties. The Pocket Negotiator should also help negotiators become aware of the role of emotions, moods, and interaction in negotiation; see, for example, Fisher, Ury, and Patton (1992) and Thompson (2005). For example, to help the user regulate emotions (his or her own or those of the opponent), the agent should have some means of establishing the emotional state of the user (and preferable that of the opponent), the agent should know the conflict-handling style of the user (and preferably that of the opponent), and the agent should be able to link emotions to core concerns (appreciation, affiliation, autonomy, status, and role) (see Fisher and Shapiro 2005).

The agent needs to incorporate general knowledge about emotions, coping styles, and mental models. Emotions or moods, for example, are triggered by a conglomerate of factors such as situation, context, interaction with other people, and physical state (see, for example, Ursin and Eriksen [2004]). Successful behavioural responses grow into coping styles (Lazarus and Folkman 1984) of that individual.

Tools and techniques are needed to elicit information from the user on the conflict-handling styles of both parties (Thomas 1992) and on the mental model of negotiation of the user (Boven and Thompson 2003). This knowledge is to form the basis of an agent that provides general coping advice that fits the profile of the user and is relevant for the situation the user is in.

Bidding Support

To properly assist the user, the Pocket Negotiator has to give run-time advice on bidding strategies, on the quality of bids received from the opponent, on possible counteroffers that the user can make, on whether to accept an offer, to walk away, or to continue with the negotiation. Essential in this process is giving the user insight into the bidding history and a prognosis of future developments (see for example, Kersten and Gray [1996]). An idea for bidding support is illustrated in figure 1, where the user is presented with the space of possible bids plotted on the basis of the utility of the user and the estimated utility of the negotiation partner. By pointing to a bid in the space, the interface presents the details of that bid on screen. Fundamental questions underlying these issues refer to the research into computationally efficient bidding strategies that lead to win-win outcomes and cannot be exploited by the opponent (see for example,

Jonker and Treur [2001]; Ludwig, Kersten, and Huang [2006]); the research in this area is ongoing. Also techniques must be improved to reduce the complexity of the negotiation space while maintaining accuracy in bidding (Hindriks, Jonker, and Tykhonov 2006). Heuristics must be developed for run-time estimation of the Pareto-efficient frontier and efficient outcomes, such as Nash, Kalai-Smorodinski (Raiffa, Richardson, and Metcalfe 2002). So far, the computational complexity of these questions has not been tackled. There is a great need for research and development of evaluation tools and techniques for the analysis of the dynamics of negotiation (Bosse and Jonker 2005; Hindriks, Jonker, and Tykhonov 2007; Jonker and Treur 2001; Kersten and Cray 1996). Through on-screen visualisation the Pocket Negotiator enhances the user's awareness of the negotiation space, potential strategies, and the interests of the opponent (Spence 2007). Many questions remain in this area. Especially the relation between the bidding process and the negotiation outcome still remains unclear. Tools and techniques must be created to assist the professional user in selecting an appropriate bidding heuristic and to fine-tune that heuristic.

We believe it is particularly interesting to develop support that can work with incomplete and qualitative information. Research is needed to clarify the relation between qualitative representations of the preferences and other information about the domain being negotiated, that is, the belief state of a negotiator. This is an important area of research as it may help clarify when to make what type of negotiation move, that is, when to provide an offer, to ask a question, or provide information to an opponent.

Learning Opponent Preferences

To reach an agreement in bilateral negotiation both parties aim to satisfy their own interests. However, to reach an agreement at all, they have to take their opponent's preferences into account. Negotiating parties generally are unwilling to reveal their preferences in order to avoid exploitation. As a result, both parties have incomplete information, which makes it hard to decide on a good negotiation move and hard to reach an optimal agreement.

Even though software agents can outperform humans in well-defined negotiation domains (see for example, Bosse and Jonker [2005]), in general such agents cannot reach optimal outcomes either without sufficient knowledge about the negotiation domain or their opponents. As negotiation is recognized as an important means for agents to achieve their own goals efficiently (Rosenschein and Zlotkin 1994) the challenge thus is to maximize the performance of automated negotiation

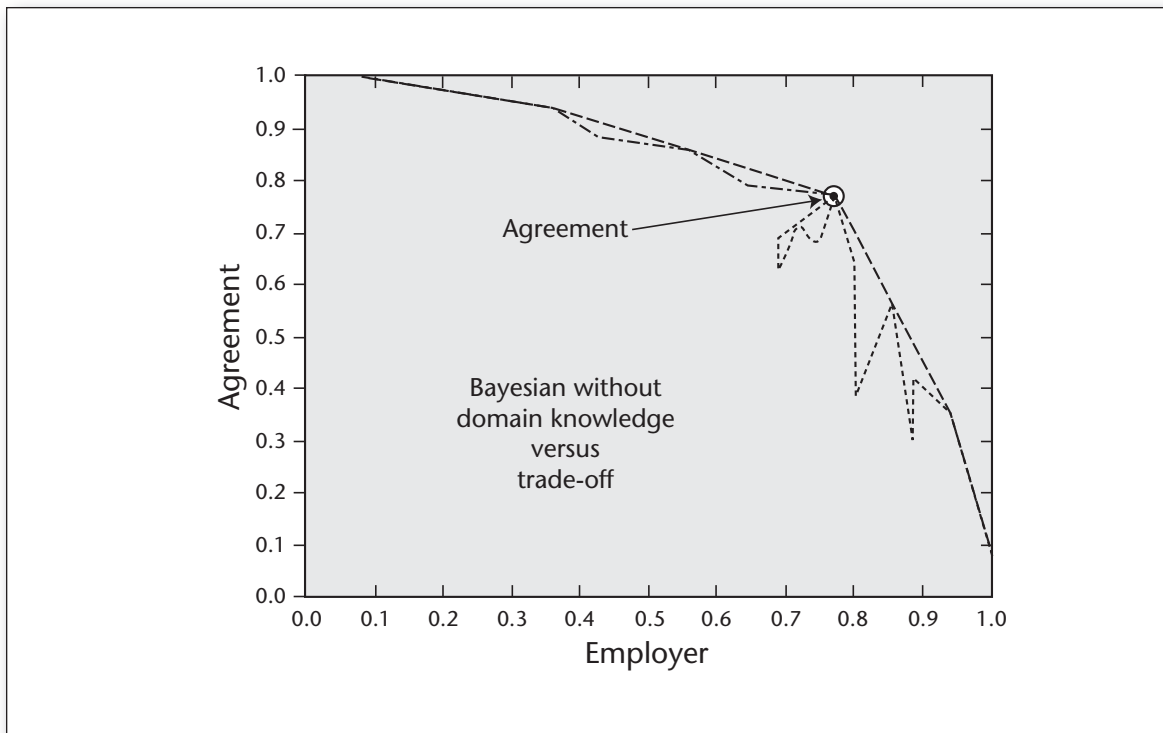


Figure 1. Employee-Employer Negotiation Domain.

agents given this limited availability of information. Research analyzing various negotiation domains and algorithms (see, for example, Hindriks, Jonker, and Tykhonov [2007]; Faratin, Sierra, and Jennings [2003]; Zeng and Sycara [1997]) has shown that efficient negotiation requires both knowledge about the negotiation domain as well as about opponent preferences.

However, the private preferences of an agent will not simply be revealed to an opponent. Generally it is unwise to reveal information about what is minimally acceptable (your reservation price) since this will provide an opponent with the opportunity to force this outcome (Raiffa 1982). If the negotiating parties have a sufficient amount of trust in each other, some information might be volunteered. Humans might also offer feedback about the bids received from the opponent (for example, your last bid is actually worse than your previous bid). If no information is offered freely, an alternative to obtain information about an opponent's private preferences is to derive it from the negotiation moves performed by that opponent during a negotiation. Various learning techniques have been proposed to uncover private preferences (Coehoorn and Jennings 2004; Faratin, Sierra, and Jennings 2003; Jonker, Robu, and Treur 2007; Zeng and Sycara 1998; Hindriks and Tykhonov 2008). A complicating factor is that the number of moves performed before reaching an agreement is limited (typically about 5 to 30 moves), and individual

bids do not provide much information (Zeng and Sycara 1997).

It is nonetheless possible to construct an opponent model, that is, a model of the opponent's preferences that can be effectively used to improve negotiation outcomes. The main idea is to exploit certain structural features and rationality principles to limit the possible set of preference profiles that can be learned. We present a learning algorithm based on Bayesian learning techniques that uses assumptions about the structure of opponent preferences and the rationality of the bidding process itself. This approach can be integrated into various negotiation strategies since the main focus is on learning an opponent's utility space.

In order to ensure that learning in a single negotiation encounter with another negotiating agent is feasible, it is essential to make some reasonable assumptions. The first assumption is a common one, (see for example, Raiffa [1982]), and assumes that the utility of a bid can be computed as a weighted sum of the utilities associated with the values for each issue. Utility functions modeling the preferences of an agent thus are linearly additive functions. In order to learn an opponent's preference profile or utility function we thus need to learn both the issue priorities or weights w_i as well as the evaluation functions $e_i(x_i)$. The objective of learning an opponent model now is to find a model that is the most plausible candidate or best

approximation of the opponent's preference profile.

The idea is to learn an opponent preference profile from its negotiation moves, that is, the bids it proposes during a negotiation. In a Bayesian learning approach, this means we need to be able to update the probability associated with all hypotheses given new evidence, that is, one of the bids. More precisely, we want to compute $P(h_j|b_t)$ where b_t is the bid proposed at time t . In order to be able to use Bayes's rule to do this, however, we need some information about the utility the opponent associates with bid b_t .

As this information is not generally available, we introduce an additional assumption to be able to make an educated guess of the utility value of b_t for an opponent. The assumption that we use is that our opponent follows a more or less rational strategy in proposing bids. In particular, we will assume that an opponent follows some kind of concession-based strategy. Although assuming such behaviour may not always be realistic it typically is necessary to perform at least some concession steps in order to reach an agreement. Moreover, in game-theoretic approaches and in negotiation it is commonly assumed that agents use a concession-based strategy (Faratin, Sierra, and Jennings 1998; Osborne and Rubinstein 1994).

In line with Faratin, Sierra, and Jennings (1998) we assume that a rational agent starts with a bid of maximal utility and moves according to a monotonically decreasing function toward its reservation value when approaching the negotiation deadline. This assumption still allows that an opponent uses various kinds of tactics and no exact knowledge about an opponent's negotiation tactics is assumed. More specifically, the rationality assumption is modeled as a probability distribution associated with a range of tactics. As a result, each utility associated with an opponent's bid also has an associated probability.

This assumption allows us to compute a predicted utility value $u'(b_t)$ for an opponent's bid b_t , which in turn allows us to compute the conditional probability $P(b_t|h_j)$ representing the probability of bid b_t given hypothesis h_j at time t . This is done by defining the probability distribution $P(b_t|h_j)$ over the predicted utility of b_t using the rationality assumption and the utility of b_t according to hypothesis h_j . Here the predicted utility $u'(b_t)$ of a next bid of the opponent is estimated as $u'(b_{t-1}) - c(t)$ using a function $c(t)$ that is the most plausible model of the negotiation concession tactic used by the opponent.

Figure 1 shows the results of the experiments, including the negotiation traces and the Pareto-efficient frontier. The agreement reached is also marked explicitly. In the domain used in figure 1, the setting is that of an employee and an employ-

er who negotiate about a job assignment and related issues such as salary. An interesting aspect of this domain is that both parties have the same preferences with regards to one of the issues. The Bayesian agents are able to reach an agreement close to the Pareto-efficient frontier. More information on this learning method can be found in Hindriks and Tykhonov (2008).

Negotiation Strategy

Two basic, constitutive facts about negotiation define the basic dilemma each negotiator has to face: (1) each party aims to satisfy its own interests as best as possible, but (2) in order to reach an agreement one has to take one's opponent's preferences into account as well.

In the literature on automated negotiation, typically, concession-based strategies have been proposed. An agent that uses a concession-based strategy selects as the next offer it will make an offer that has a decreased utility compared with the last offer made. The utility that is being decreased is the utility from the agent's own perspective without any guarantee that such a decrease will also increase the utility from the other party's perspective. A well-known example of such a strategy is the time-dependent strategy, which decreases utility simply as a function of time (Faratin, Sierra, and Jennings 1998).

Although motivated by fact 2 above, such strategies do not explicitly take the opponent's preferences into account, and, as a result, will most likely be inefficient in complex negotiation domains. Moreover, time-dependent strategies can be exploited by the other negotiating party and as such do not adequately take fact 1 above into account.

The solution to these problems is to explicitly take the preferences of an opponent into account. One key question still needs to be addressed: How can an agent exploit information about opponent preferences effectively?

The preferences of an opponent can be used in at least two ways. First, it can be used to propose efficient Pareto-optimal offers. Finding such offers requires that the Pareto frontier can be approximated, which is feasible only if a reasonable model of the opponent's preferences is available. Second, it can be used to recognize and avoid exploitation. The strategy we discuss is inspired by a classification of negotiation moves as described in Hindriks, Jonker, and Tykhonov (2007) and the Tit-for-Tat tactic, discussed in Axelrod (1984) and — in a negotiation context — in Faratin, Sierra, and Jennings (1998).

The main criteria are that the strategy should be efficient and transparent, maximize the chance of an agreement, and avoid exploitation.

The first observation relevant to the design of

the strategy is that the availability of information about the preferences of an opponent enables an agent to classify the moves its opponent makes. Here, we use a classification of moves proposed in Hindriks, Jonker, and Tykhonov (2007) and illustrated in figure 2. The move classification is presented from the perspective of agent A.

Given that agent A's last offer is marked by the arrow "Current Bid of Agent A," the agent has a number of choices for making a next negotiation move. A silent move does not change utility of either party significantly. A concession move decreases own utility but increases the utility of the opponent. A fortunate move increases utility for both parties, whereas an unfortunate move does the opposite. A fortunate move can only be made if the current bid is not already on the Pareto frontier. A selfish move increases own utility but decreases the opponent's utility. Finally, a nice move increases the opponent's utility but does not change the agent's own utility.

Based on this classification a suggestion would be to "mirror" each move of an opponent by making a similar move, which would implement a Tit-for-Tat-like tactic. The basic idea of a Tit-for-Tat strategy in a multiissue negotiation context would be to respond to an opponent move with a symmetrical one. That is, "match" the move as depicted in figure 3 by mirroring it in the diagonal axis.

First note that each type of move would indeed result in a response move in the same class. In particular, responding to a concession move of the opponent with a concession move itself arguably is one of the most reasonable responses one can make. All rational negotiation strategies will attempt to make concession moves at some point during a negotiation. Moreover, the "mirroring" strategy would avoid exploitation as a selfish move of the opponent would result in a selfish response move. Such a response would be a signal to the opponent, "I am prepared to make a concession toward you only if I get something in return. If you pull back I'll do the same."

A mirroring strategy would, however, be too simplistic for several reasons. A mirroring strategy is not rational in the case of an unfortunate move, as there is no reason to decrease the agent's own utility without increasing the chance of acceptance of the proposed bid by the opponent. Furthermore, observe (compare figure 2) that unfortunate moves move away from the Pareto-optimal frontier, and thus would not satisfy our efficiency criteria.

In order to remove these deficiencies, we propose first to mirror the move of the opponent and thereafter make an additional move toward the Pareto frontier, that is, a move toward the approximated Pareto frontier that is computed using the learned opponent model and the agent's own preference profile. There are multiple ways to do this

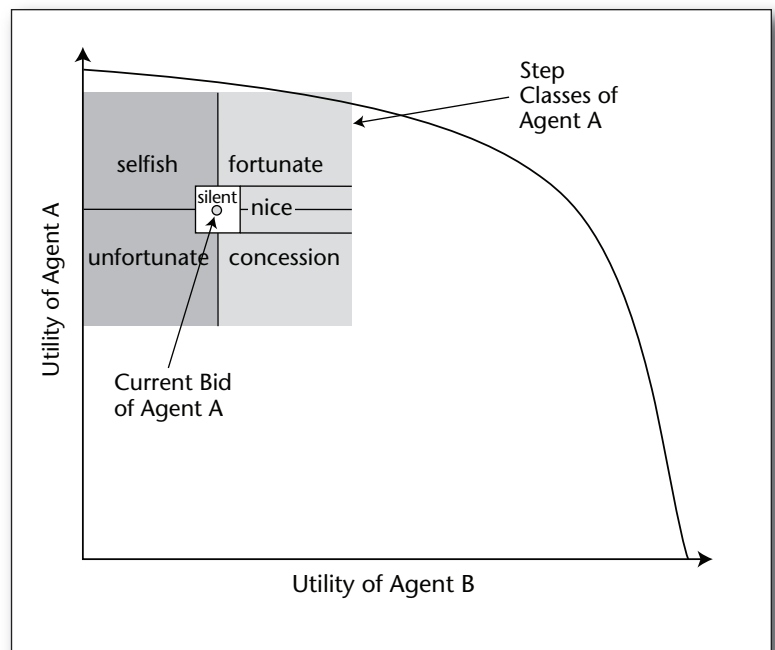


Figure 2. Classification of Negotiation Moves

and the choice is not straightforward. What is clear is that the move toward the Pareto frontier should not further decrease the agent's own utility as this would invite exploitation tactics. It also does not seem rational to further decrease the opponent's utility as this would result in selfish moves to arbitrary moves of the opponent.

The final observation that motivates this choice is that increasing the agent's own utility by moving toward the Pareto frontier actually minimizes the chance of reaching an agreement when this strategy would be used by both parties, which would violate one of our design criteria for a negotiation strategy. To explain this, consider two agents that would mirror an opponent's move and then, seen from the perspective of Agent A in figure 3, would move straight up toward the Pareto frontier (Agent B would move right), which would only increase own utility. The other agent in this case would consider such a move a selfish move and respond similarly, thereby minimizing the chance of reaching an agreement. Of course, this line of reasoning depends on the quality of the opponent model but presents a real problem. To resolve it, the strategy we propose only increases the opponent's utility when moving toward the Pareto frontier in order to maximize the chance of an agreement. The resulting strategy consists of two steps: first mirror the move of the opponent and then add a nice move to propose an efficient offer (that is, search for a bid on the approximated Pareto frontier that is on the same iso-curve as the bid obtained by mirroring; see figure 3). This strat-

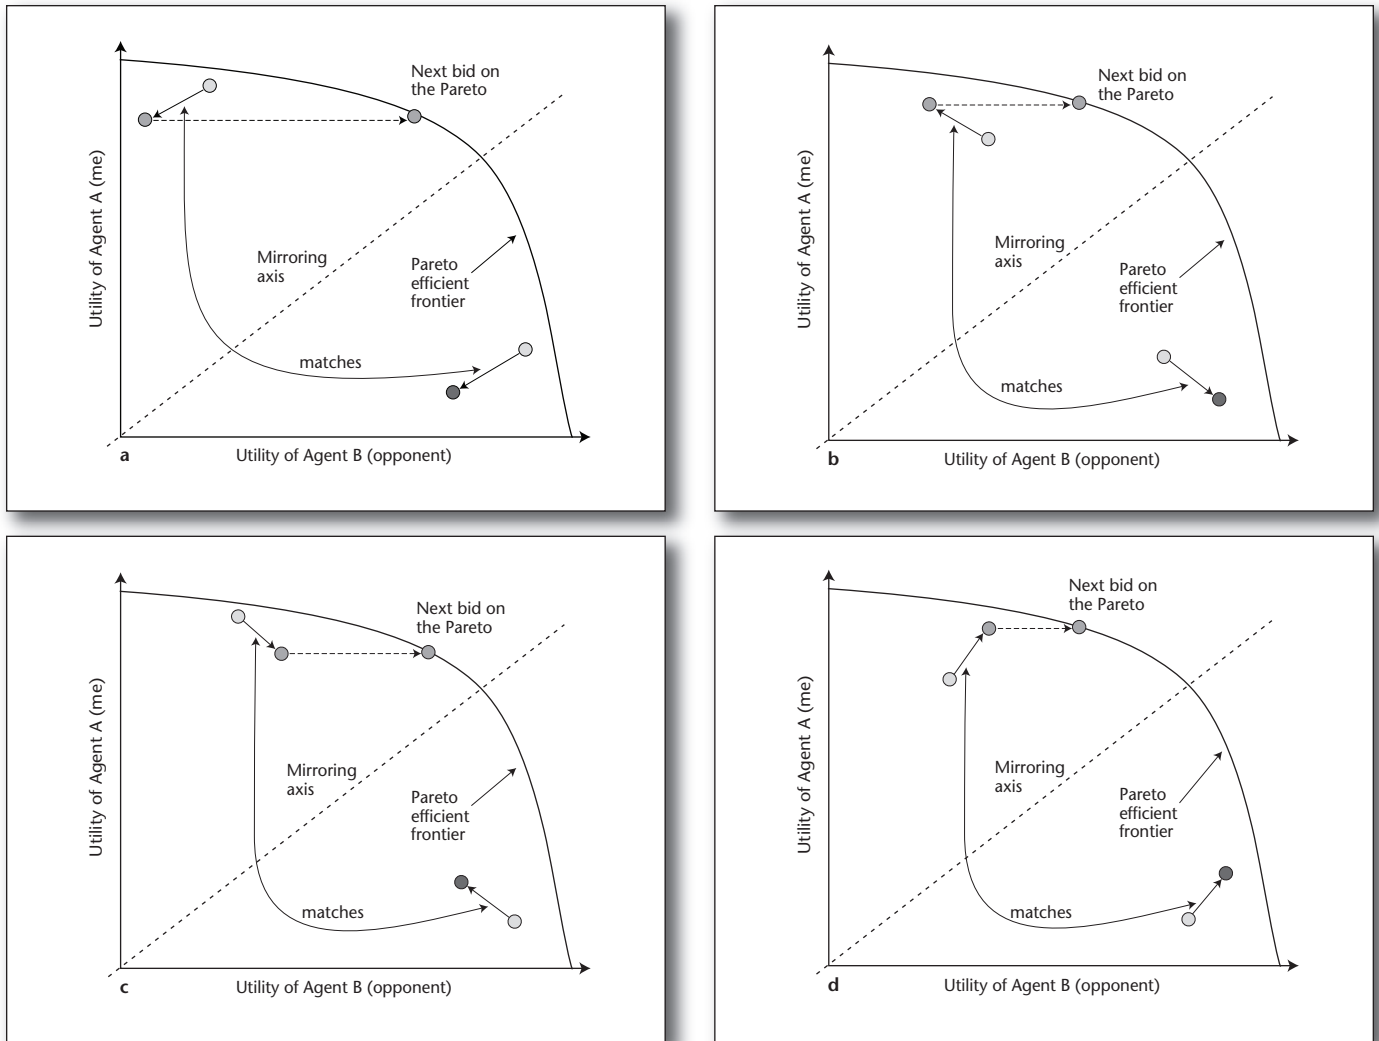


Figure 3. Example Responses.

(a) Unfortunate move, (b) selfish move, (c) concession move, and (d) fortunate move.

egy we call the *mirroring strategy* (MS). To gain a better understanding of MS, it is instructive to discuss some of the response moves MS generates. Figure 3 shows examples of responses to an unfortunate, selfish concession and fortunate move. The response to an unfortunate move is to mirror this move and add a nice move, which results in a concession move (see figure 3a). This is a reasonable reply, as such a move may be interpreted as an attempt (that failed) to make a concession move by the opponent (due to the lack of information about the preferences of its opponent). Such a move, which is the result of misinformation, should not be punished, we believe, but an attempt instead should be made to maintain progress toward an agreement.

The response to a selfish move results either in a fortunate move or in a selfish move. Figure 3b shows the case resulting in a fortunate move. It should be noted that a fortunate move is only pos-

sible if the previous move the agent made was inefficient. This means that in that case the opponent model must have misrepresented the actual preferences of the opponent. In such a case, where our previous move was based on misinformation, we believe it is reasonable to not punish the opponent with a selfish move and give the opponent the benefit of the doubt in such a case. If, however, the previous move would have been efficient, a selfish move most likely would be replied to with a selfish move (since there would be no room to make a nice move toward the Pareto frontier), and it is reasonable to send a clear signal to the opponent that such moves are unacceptable.

Finally, both a concession move as well as an unfortunate move of the opponent would be replied to with the same type of move (see figures 3c and 3d). Moreover, if there is room for a nice move toward the Pareto frontier, in both cases the step would be bigger than that made by the oppo-

ment, increasing the utility of the opponent even more and thereby again increasing the chance of acceptance as early on in a negotiation as possible.

As discussed, a negotiation strategy should be efficient and transparent, maximize the chance of an agreement, and avoid exploitation. It is clear that MS aims to be as efficient as possible, and this depends on the quality of the learning technique for modeling opponent preferences. MS is transparent as it proposes a simple response strategy by mirroring an opponent's move and then adding a nice step. The signals thus sent by negotiation moves are easy to interpret by an opponent. In particular, MS only punishes an opponent in reply to a selfish move and does so only when the model of opponent preferences matches the actual preferences of that opponent. As a result, MS not only avoids exploitation but also is a nice strategy. MS is nice even when an opponent makes unfortunate moves that are interpreted as "mistakes" on the opponent's part. The strategy moreover maximizes the chance of an agreement as early as possible, which is achieved by the move toward the Pareto frontier that always maximizes the utility of the opponent relative to a particular utility for the agent itself.

Human Preference Elicitation

In negotiation support, the quality of the outcome depends to a large extent on the quality of the preparation of the negotiators and their interaction. Both preparation and interaction should focus on discovering the preferences of both parties (Fisher, Ury, and Patton 1992).

Eliciting preferences is not simple. Confronted with a new decision task, people do not possess known, stable, and coherent preferences. Instead, preferences are constructed at the time the valuation question is asked (Payne, Bettman, and Schkade 1999a; Simon 1955; Curhan, Neale, and Ross 2004). Furthermore, the decision process itself and the context play a major role in the construction process (Payne, Bettman, and Johnson 1999b). This includes the alternatives in an outcome set and how information is presented to the person.

Given that the process of constructing preferences is important for people to arrive at an understanding of their own preferences and as the flow of the process influences the outcome, it is important to design that process carefully, so that the user is able to construct an accurate model. We believe that a major factor in the process is the interaction between the system and its user through a preference elicitation interface. Therefore, in order to create more successful systems that can elicit accurate preferences we have to focus on the design of the user interface. Even the best underlying algorithms and reasoning frameworks

do not give successful results if the user has problems interpreting information presented by the system and entering his or her preferences (Peintner and Paolo Viappiani. 2008).

By actively involving the participants in the design process we were able to understand how they prefer an interface to be designed. We learned that an important aspect of the process design is that it allows people to understand their own preferences and that people feel in charge of creating their profile as opposed to just answering questions that are used by the system to build the profile. In particular, being able to explore their preferences from different angles including underlying interests and consequences (in form of rankings of decision outcomes) supported people's process of constructing their preferences. Participants like design elements that support this exploration in a natural way and provide immediate visual feedback. This is consistent with design guidelines established earlier by Pu and Chen (2008).

Based on the results of this study (Pommeranz, Wiggers, and Jonker 2010) and a study on different ways of entering preferences, we established the four following design guidelines for preference elicitation interfaces:

- (1) As motivated users are willing to spend more effort, users should be given the option to express more detail if they feel the need to do so.
- (2) As affective feedback was the preferred way of adding finer-grained preference detail, interfaces should consider affective feedback as a useful mechanism for specifying detailed preference feedback.
- (3) The user must be able to explore his/her interests, preferences, and outcomes in the same physical space in a way that gives immediate feedback on the links between the three concepts.
- (4) The user's cognitive load of giving preferences can be reduced by showing default preferences based on profile/interest selection that can subsequently be adapted.

These guidelines and the data collected in the evaluations informed our further design process (Pommeranz, Wiggers, and Jonker 2010, 2011). The preference elicitation interface of our current prototype is shown in figure 4. The interface has three panels: (1) one where people can specify their interests, (2) one for entering preferences using post-it notes, and (3) one that shows relevant offers. The interests panel is inspired by the value-focused thinking approach (Keeney 1992). The idea is that people find it easier to specify their values than their specific preference for a particular domain, as values are typically more stable during one's life. Based on these interests the systems fills in a number of preference suggestions in the second panel.

When a user changes his or her preferences, the contents of the third panel, which shows hypo-

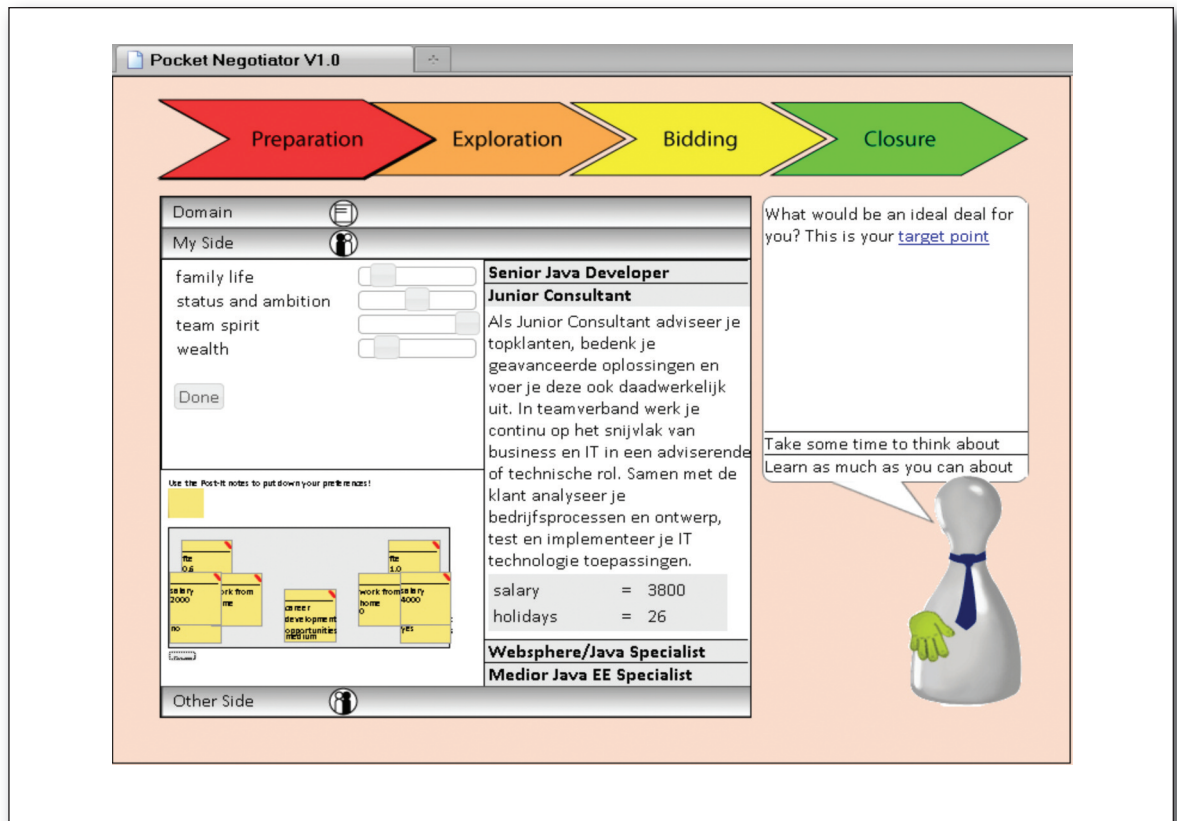


Figure 4. Pocket Negotiator Prototype — Preference Elicitation Interface.

thetical offers that fit the current preference profile, change as well. This immediate visual feedback helps users to understand the consequences of their choices and to refine their preference profile.

Modeling Emotion in Negotiation

Emotion and affect play an important role in human cognition and behavior. This is no different for the role of affect in negotiation (for a review see Broekens, Jonker, and Meyer [2010]). For example, the expression of anger communicated by one negotiation partner can influence the other partner to concede more, but only if the other partner is in a low-power (low control) situation (Kleef et al. 2006). Strong negative feelings and especially a bad atmosphere almost always hinder the negotiation process (Luecke 2003), while a mildly negative mood can favor critical thinking.

Positive moods favor creative thought (Isen, Daubman, and Nowicki 1987), needed to create value in negotiations, while enthusiasm can result in being not critical enough. Finally, a person's belief about something is updated according to the emotions of that person: the current emotion is used as information about the perceived object (Gasper and Clore 2002; Forgas 2000), and emo-

tion is used to make the belief resistant to change (Frijda and Mesquita 2000). This is important for preferences and preference elicitation.

An NSS that takes affect into account can (1) help the user of the system be aware of his or her own emotion, mood, and preferences (for example, if you seem to feel sad, be aware of the fact that this makes you feel pessimistic about the negotiation); (2) organize the negotiation process to be more compatible with the user's mood, or try to influence the user's mood to be more compatible with the current negotiation activities (for example, creative thinking is important in the joint exploration phase to come up with as much value as possible, so a user in this phase could be primed to be in a positive mood); (3) indicate to the user when it is opportune to make strategic use of emotions (for example, express anger to claim value now); (4) detect the expressed emotion of the negotiation partner and help the user analyze the meaning of this signal (for example, your partner shows anger, but is not in a dominant position; you should try to ignore the expression and keep being constructive); (5) provide offline training by means of virtual reality role play including affective intelligent virtual agents (Core et al. 2006; Traum et al. 2008) (for example, play and reflect

upon a virtual reality negotiation scenario).

The affective support functions of an NSS are based upon two different pillars (Broekens, Jonker, and Meyer 2010). The first is knowledge about relations between affect and negotiation, preferably grounded in actual experimental studies. This knowledge defines the kind of advice the affective NSS should give. For example, positive moods favor creative thought, so joint exploration should be done in a positive mood, while slightly negative moods favor critical thinking and attention to details, so perhaps bidding should be done in a more neutral to negative mood. The second pillar is affective computing and affective human-machine interaction methods and techniques (Pickard 1997, Hudlicka 2003). These methods and techniques define what is possible to do. For example, the ability to make the user aware of its mood depends on a method that can measure that mood, while the ability to interpret the affective expression of the user's negotiation partner depends on a method to detect affective expressions.

Conclusion

The creation of negotiating agents that can support humans with their negotiations is a multidisciplinary challenge. The inherent semantic problem and the emotional issues involved mean that negotiation cannot be handled by artificial intelligence alone, and a human-agent team is required. By performing an analysis on the potential strengths and weaknesses of the team members, key areas of agent technology, artificial intelligence, and human-computer interaction are identified that need to be addressed to develop such a team: automated negotiating agents such as these outperform human beings in terms of deal optimality, with affecting computing for handling emotions, and preference elicitation to know what is important in the negotiation. This article provides references to advances in associated research areas and describes key results of the team working on the Pocket Negotiator agent.

Areas scarcely or not addressed in this article that are important for the Pocket Negotiator agent are shared mental models and team models, explanation, argumentation, value elicitation, value sensitive design, culture dependence, mediation, and multiparty negotiation.

Acknowledgements

This research is supported by the Dutch Technology Foundation STW, applied science division of NWO, and the Technology Program of the Ministry of Economic Affairs. It is part of the Pocket Negotiator project with grant number VICI-project 08075.

Notes

1. See interneg.carleton.ca/inspire.
2. See www.athenasoft.org.
3. See www.smartsettle.com.
4. See the publications on the agents of ANAC 2010 and those of ANAC 2011 as published in the proceedings of the AAMAS workshop ACAN.
5. This system is available for free download from mmi.tudelft.nl/negotiation/index.php/Genius.

References

- Axelrod, R. 1984. *The Evolution of Cooperation*. New York: Basic Books, Inc.
- Bosse, T., and Jonker, C. M. 2005. Human Versus Computer Behaviour in Multiissue Negotiation. In *Proceedings of the First International Workshop on Rational, Robust, and Secure Negotiations in Multiagent Systems*, ed. T. Ito, H. Hattori, T. Matsuo, and M. Zhang, 10–25. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Boven, L. van, and Thompson, L. 2003. A Look into the Mind of the Negotiator: Mental Models in Negotiation. *Group Processes and Intergroup Relations* 6(4): 387–404.
- Broekens, J.; Jonker, C. M.; and Meyer, J.-J. C. 2010. Affective Negotiation Support Systems. *Journal of Ambient Intelligent Smart Environments* 2(2): 121–144.
- Brazier, F. M. T.; Jonker, C. M.; Treur, J.; and Wijngaards, N. J. E. 2000. On the Use of Shared Task Models in Knowledge Acquisition, Strategic User Interaction and Clarification Agents. *International Journal of Human-Computer Studies* 52(1): 77–110.
- Büttner, R. 2006. The State of the Art in Automated Negotiation Models of the Behavior and Information Perspective. *International Transactions on Systems Science and Applications* 1(4): 351–356.
- Chen, L. S. 2000. Joint Processing of Audio-Visual Information for the Recognition of Emotional Expressions in Human-Computer Interaction. Ph.D. thesis, University of Illinois at Urbana-Champaign, Department of Electrical Engineering, Urbana, IL.
- Coehoorn, R. M., and Jennings, N. R.; 2004. Learning an Opponent's Preferences to Make Effective Multiissue Negotiation Trade-Offs. In *Proceedings of 6th International Conference on Electronic Commerce*, 59–68. New York: Association for Computing Machinery.
- Core, M.; Traum, D.; Lane, H. C.; Swartout, W.; Gratch, J.; van Lent, M.; and Marsella, S. 2006. Teaching Negotiation Skills Through Practice and Reflection with Virtual Humans. *Simulation* 82(11): 685–701.
- Curhan, J. R.; Neale, M. A.; and Ross, L. 2004. Dynamic Valuation: Preference Changes in the Context of Face-to-Face Negotiation. *Journal of Experimental Social Psychology* 40(2): 142–151.
- DeLuca, M. J., and DeLuca, N. F. 2007. *Negotiating Salary and Job Offers. Hundreds of Ready-to-Use Phrases to Help You Get the Best Possible Salary, Perks, or Promotion*. New York: McGraw-Hill.
- Faratin, P.; Sierra, C.; and Jennings, N.R. 2003. Using Similarity Criteria to Make Negotiation Trade-Offs. *Journal of Artificial Intelligence* 142(2): 205–237.
- Faratin, P.; Sierra, C.; and Jennings, N. R. 1998. Negotia-

- tion Decision Functions for Autonomous Agents. *International Journal of Robotics and Autonomous Systems* 24(3–4): 159–182.
- Fisher, R., and Shapiro, D. 2005. *Beyond Reason: Using Emotions as You Negotiate*. New York: Random House Business Books.
- Fisher, R.; Ury, W. L.; and Patton, B. 1992. *Getting to Yes: Negotiating Agreement Without Giving In*. London: Penguin Books.
- Folkman, S., and Lazarus, R. S. 1990. Coping and Emotion. In *Psychological and Biological Approaches to Emotion*, ed. N. L. Stein, B. Leventhal, and T. Trabasso, 313–332. Hillsdale, NJ: Lawrence Erlbaum.
- Forgas, J. P. 2000. Feeling Is Believing? The Role of Processing Strategies in Mediating Affective Influences in Beliefs. In *Emotions and Beliefs: How Feelings Influence Thought*, ed. N. H. Frijda, A. S. R. Manstead, and S. Bem, 108–143. Cambridge, UK: Cambridge University Press.
- Frijda, N. H., and Mesquita, B. 2000. Beliefs Through Emotions. In *Emotions and Beliefs: How Feelings Influence Thought*, ed. N. H. Frijda, A. S. R. Manstead, and S. Bem, 45–77. Cambridge, UK: Cambridge University Press.
- Gasper, K., and Clore, G. L. 2002. Attending to the Big Picture: Mood and Global Versus Local Processing of Visual Information. *Psychological Science* 13(1): 34–40.
- Hindriks, K. V.; Jonker, C. M.; and Tykhonov, D. 2006. Reducing Complexity of an Agent's Utility Space for Negotiating Interdependent Issues. Paper presented at the International Conference on Complex Systems, Boston, June 25–30.
- Hindriks, K. V.; Jonker, C. M.; and Tykhonov, D. 2007. Negotiation Dynamics: Analysis, Concession Tactics, and Outcomes. In *Proceedings of the IEEE/WIC/ACM Conference on Intelligent Agent Technologies*, 427–433. Los Alamitos, CA: IEEE Computer Society.
- Hindriks, K. V.; Jonker, C. M.; and Tykhonov, D. 2007. Analysis of Negotiation Dynamics. In *Proceedings of 11th International Workshop on Cooperative Information Agents*, ed. M. Klusch, K. V. Hindriks, M. P. Papazoglou, and L. Sterling, 27–35. Berlin: Springer.
- Hindriks, K. V., and Tykhonov, D. 2008. Opponent Modelling in Automated Multiissue Negotiation. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, 331–338. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Hofstede, G. J.; Jonker, C. M.; and Verwaart, T. 2010. Cultural Differentiation of Negotiating Agents. *Group Decision Negotiation* 21(1): 1–20.
- Hudlicka, E. 2003. To Feel or Not to Feel: The Role of Affect in Human-Computer Interaction. *International Journal of Human-Computer Studies* 59(1–2): 1–32.
- Isen, A. M.; Daubman, K. A.; and Nowicki, G. P. 1987. Positive Affect Facilitates Creative Problem Solving. *Journal of Personality and Social Psychology* 52(6): 1122–1131.
- Ito, T.; Zhang, M.; Robu, V.; Fatima, S.; Matsuo, T.; Yamaki, H. (eds.) 2010. *Innovations in Agent-Based Complex Automated Negotiations*, Volume 319 of Series on Studies in Computational Intelligence, 228. Berlin: Springer-Verlag.
- Jonker, C. M., and Treur, J. 2001. An Agent Architecture for Multiattribute Negotiation. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, 1195–1201. San Francisco: Morgan Kaufmann Publishers.
- Jonker, C. M.; Robu, V.; and Treur, J. 2007. An Agent Architecture for Multiattribute Negotiation Using Incomplete Preference Information. *Journal of Autonomous Agents and Multiagent Systems* 15(2): 221–252.
- Jonker, C. M.; Riemsdijk, M. B. van; and Vermeulen, B. 2011. Shared Mental Models: A Conceptual Analysis. In *COIN 2010 International Workshop*, Lecture Notes in Artificial Intelligence Volume 6541, 32–151. Berlin: Springer.
- Keeney, R. 1992. *Value-Focused Thinking: A Path to Creative Decision Making*. Cambridge, MA: Harvard University Press.
- Kersten, G. E., and Cray, D. 1996. Perspectives on Representation and Analysis of Negotiation: Towards Cognitive Support Systems. *Group Decision and Negotiation* 5(4–6): 433–467.
- Kleef, G. A. V.; De Dreu, C. K. W.; Pietroni, D.; and Manstead, A. S. R. 2006. Power and Emotion in Negotiation: Power Moderates the Interpersonal Effects of Anger and Happiness on Concession Making. *European Journal of Social Psychology* 36(4): 557–581.
- Kraus, S.; 2001. *Strategic Negotiation in Multiagent Environments*. Cambridge, MA: The MIT Press.
- Lazarus, R. S., and Folkman, S. 1984. *Stress, Appraisal, and Coping*. Berlin: Springer.
- Ludwig, S. A.; Kersten, G. E.; and Huang, X. 2006. Towards a Behavioural Agent-Based Assistant for e-Negotiations. Paper presented at the Montreal Conference on e-Technologies (MCETECH). Montreal, Quebec, Canada, May 17–19.
- Luecke, R. 2003. *Harvard Business Essentials: Negotiation*. Cambridge, MA: Harvard Business School Press.
- Meyer, T.; Foo, N. Y.; Kwok, R.; and Zhang, D.; 2004. Logical Foundations of Negotiation: Outcome, Concession, and Adaptation. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, 293–298. Menlo Park, CA: AAAI Press.
- Oliver, J. R.; 2005. On Learning Negotiation Strategies by Artificial Adaptive Agents in Environments of Incomplete Information. In *Formal Modelling in Electroning Commerce, Handbook on Information Systems, Part IV*, 445–461. Berlin: Springer-Verlag.
- Osborne, M. J., and Rubinstein, A. 1994. *A Course in Game Theory*. Cambridge, MA: The MIT Press.
- Payne, J. W.; Bettman, J. R.; and Schkade, D. A. 1999. Measuring Constructed Preferences: Towards a Building Code. *Journal of Risk and Uncertainty* 19(1–3): 243–270.
- Payne, J. W.; Bettman, J. R.; and Johnson, E. J. 1999. *The Adaptive Decision Maker*. Cambridge, UK: Cambridge University Press.
- Peintner, B.; and Paolo Viappiani, N.-S.; 2009. Preferences in Interactive Systems: Technical Challenges and Case Studies. *AI Magazine* 29(4): 93–103.
- Picard, R. W. 1997. *Affective Computing*. Cambridge, MA: The MIT Press.
- Pommeranz, A.; Wiggers, A.; and Jonker, C. M. 2010. User-Centered Design of Preference Elicitation Interfaces for Decision Support, In *HCI in Work and Learning, Life and Leisure*, Lecture Notes in Computer Science, Volume 6389, 14–33. Berlin: Springer.
- Pommeranz, A.; Wiggers, P.; and Jonker, C. M.; 2011.

Towards Compositional Design and Evaluation of Preference Elicitation Interfaces. In *Human Centered Design*, Lecture Notes in Computer Science Volume 6776, 586–596. Berlin: Springer.

Pu, P.; and Chen, L.; 2008. User-Involved Preference Elicitation for Product Search and Recommender Systems *AI Magazine* 29(4): 93–103.

Rahwan, I.; Sonenberg, L.; and McBurney, P.; 2005. Bargaining and Argument-Based Negotiation: Some Preliminary Comparisons. In *Argumentation in Multiagent Systems*, ed. I. Rahwan, P. Moraitis, and C. Reed, 176–191. Berlin: Springer-Verlag.

Raiffa, H. 1982. *The Art and Science of Negotiation*. Cambridge, MA: Harvard University Press.

Raiffa, H.; Richardson, J.; and Metcalfe, D. 2002. *Negotiation Analysis*. Cambridge, MA: Harvard University Press.

Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. Cambridge, MA: The MIT Press.

Simon, H. A. 1955. A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics* LXIX (February): 99–118.

Spence, R. 2007. *Information Visualization — Design for Interaction*. Engelwood Cliffs, NJ: Pearson Prentice Hall.

Tamma, V.; Phelps, S.; Dickinson, I.; and Wooldridge, M. 2005. Ontologies for Supporting Negotiation in E-Commerce. *Engineering Applications of Artificial Intelligence* 18(2): 223–236.

Thomas, K. W. 1992. Conflict and Conflict Management: Reflections and Update. *Journal of Organizational Behavior* 13(3): 265–274.

Thompson, L. L. 2005. *The Heart and Mind of the Negotiator*. Engelwood Cliffs, NJ: Pearson Prentice Hall.

Traum, D.; Marsella, S.; Gratch, J.; Lee, J., and Hartholt, A. 2008. Multiparty, Multiissue, Multistrategy Negotiation for Multimodal Virtual Agents: In *Proceedings of the 8th Intelligent Virtual Agents Conference*, Lecture Notes in Computer Science 5208, 117–130. Berlin: Springer.

Ursin, H., and Erisen, H. R. 2004. The Cognitive Activation Theory of Stress. *Psychoneuroendocrinology* 29(5): 567–592.

Ury, W. 2007. *The Power of a Positive No: How to Say No and Still Get to Yes*. New York: Random House.

Ury, W. 1993. *Getting Past No. Negotiating Your Way from Confrontation to Cooperation*. New York: Random House.

Walton, R., and McKersie, R. 1965. *A Behavioral Theory of Labor Negotiations*. Thousand Oaks, CA: Sage Publications.

Zeng, D., and Sycara, K. 1997. Benefits of Learning in Negotiation. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97)*, 6. Menlo Park, CA: AAAI Press.

Zeng, D., and Sycara, K.; 1998. Bayesian Learning in Negotiation. *International Journal of Human Computer Systems* 48(1): 125–141.

Catholijn Jonker is a full professor of man-machine interaction at the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology. She studied computer science and did her Ph.D. studies at Utrecht University. Her recent publications address cognitive processes and concepts such as

trust, negotiation, and the dynamics of individual agents and organisations. In Delft she works with an interdisciplinary team to engineer human experience through multimodal interaction between natural and artificial actors in a social dynamic context. In her NWO-STW VICI project, Pocket Negotiator, she develops intelligent decision-support systems for negotiation.

Koen Hindriks is an assistant professor in the Man-Machine Interaction Group at the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology. He studied computing science, and finished his Ph.D. at Utrecht University on agent programming languages. His research interests include commonsense reasoning, agent-oriented programming based on commonsense concepts like beliefs and goals, and the verification and specification of agent programs. He has designed and developed several agent programming languages, including 3APL and GOAL. He is also interested in the design and development of negotiating agents, which involves among others research on representation, strategies, and learning techniques that can be usefully applied in the development of such agents.

Pascal Wiggers is an assistant professor in the Man-Machine Interaction Group at the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology. He received his M.Sc. cum laude in technical computer science from Delft University of Technology in 2001 and his Ph.D. from the same university in 2008. His research focuses on artificial intelligence and machine-learning techniques that enable intuitive collaborative human-machine communication and multimodal, affective interaction.

Joost Broekens is a post doctoral member of the Man-Machine Interaction Group at the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology. He received the M.Sc. degree in computer science at the University of Delft, The Netherlands, in 2001. In 2007 he received his Ph.D. in computer science at the University of Leiden, The Netherlands, in the area of computational modeling of emotion in relation to learning processes. His most recent interests include reinforcement learning, affective computing, human-robot and human-computer interaction, and gaming research.