Chapter 2
Background

Abstract In this chapter, we discuss briefly the background and related work in automated negotiation. We begin with definitions of the key aspects of automated negotiation, such as the negotiation *domain*, the *protocol*, and the *preferences*. We discuss what it means for a negotiating agent to employ a *negotiation strategy* and we highlight several prime examples of existing negotiation strategies. We also discuss a number of high-level *negotiation architectures* and how they can assist in exploring the negotiation strategy space. We focus specifically on the three components we distinguish in the Chap. 1, namely the various ways in which current negotiation strategies *bid*, *learn*, and *accept*. We conclude the background chapter by describing several methodologies for evaluating and comparing negotiation strategies and components. Among our discussed evaluation methods are *performance* and *accuracy* measures, agent competitions, and analytical software to assess the outcome of the negotiation.

2.1 Introduction

Negotiation is a common and important process for making decisions and resolving conflicts. People encounter negotiation situations everywhere, from specific situations such as job negotiations and hostage crises situations [1] to more general situations such as resource and task allocation mechanisms [2–4], conflict resolution mechanisms [5, 6], and decentralized information services [7, 8].

In recent years, the fact that negotiation covers many aspects of our lives has led to an increasing focus on the design of *automated* negotiators; i.e., autonomous agents capable of negotiating with other agents in a specific environment [7, 9]. This interest has been growing since the beginning of the 1980 s with the work of early adopters such as Smith's Contract Net Protocol [4], Sycara's PERSUADER [10, 11], Robinson's OZ [12], and the work by Rosenschein [13] and Klein [14].

In this chapter, we discuss briefly the background and related work in automated negotiation. We will begin with definitions of the basic terminology used in this

T. Baarslag, *Exploring the Strategy Space of Negotiating Agents*,
Springer Theses, DOI 10.1007/978-3-319-28243-5_2

field in Sect. 2.2. In the subsequent Sect. 2.3, we discuss several prime examples of existing negotiation strategies and their architecture. In Sect. 2.4 we discuss several ways of evaluating negotiation strategies.

2.2 Terminology

The defining elements of a bilateral negotiation are depicted in Fig. 2.1. A bilateral automated negotiation concerns a negotiation between *two* agents, usually called A and B. The party that is negotiated with is also called the *partner* or *opponent*.

The *negotiation setting* consists of the *negotiation protocol*—the rules of encounter—, the negotiating agents, and the *negotiation scenario*. The negotiation takes place in a *negotiation domain*, which specifies all possible outcomes (the so-called *outcome space*). Furthermore, every agent in the scenario has a *preference profile*, which expresses the preference relations between the possible outcomes. Together, this defines the *negotiation scenario* that takes place between the agents. The negotiation scenario and protocol specify the possible *actions* an agent can perform in a given negotiation state.

2.2.1 Negotiation Domain

The *negotiation domain*—or *outcome space*—is denoted by Ω and defines the set of possible negotiation outcomes. The *domain size* is the number of possible outcomes $|\Omega|$. A negotiation domain consists of one or more *issues*, which are the main resources or considerations that need to be resolved through negotiation; for example, the price or the color of a car that is for sale. Issues are also sometimes referred

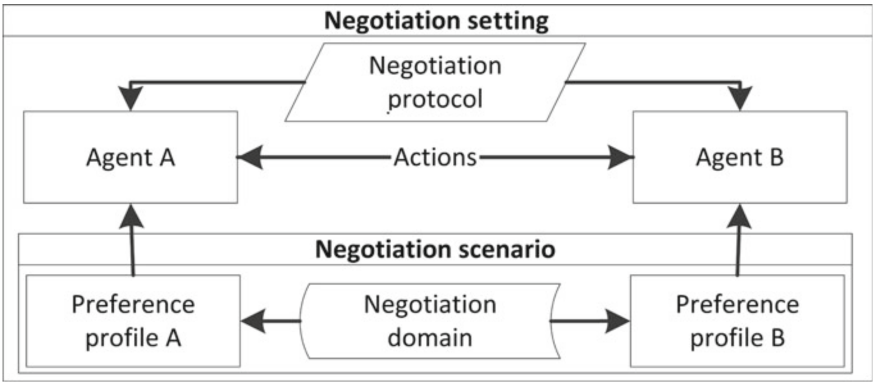

Fig. 2.1 Overview of the defining elements of an automated bilateral negotiation

to as attributes, but we reserve the latter term for *opponent attributes*, which are properties that may be useful to model to gain an advantage in a negotiation.

To reach an agreement, the agents must settle on a specific alternative or *value* for each negotiated issue. That is, an *agreement* on n issues is an outcome that is accepted by both parties of the form $\omega = \langle \omega_1, \ldots, \omega_n \rangle$, where ω_i denotes a value associated with the ith issue. We will focus mainly on settings with a finite set of discrete values per issue. A *partial agreement* is an agreement on a subset of the issues. We say that an outcome space defined by a single issue is a *single-issue negotiation*, and a *multi-issue negotiation* otherwise.

2.2.2 Negotiation Protocol

A negotiation protocol fixes the rules of encounter [15], specifying which actions each agent can perform at any given moment. Put another way, it specifies the admissible *negotiation moves*. There are a number of bilateral negotiation protocols. We do not aim to provide a complete overview of all protocols, instead we refer to Lomuscio et al. [16] for an overview of high-level parameters used to classify them, and to Marsa-Maestre et al. [17] for guidelines on how to choose the most appropriate protocol to a particular negotiation problem.

An often used negotiation protocol in bilateral automated negotiation is the *alternating offers protocol* [18, 19]. This protocol dictates that the two negotiating agents propose *outcomes*, also called *bids* or *offers*, in turns. That is, the agents create a *bidding history*: one agent proposes an offer, after which the other agent proposes a counter-offer, and this process is repeated until the negotiation is finished, for example by time running out, or by one of the parties accepting.

We use the alternating offers protocol throughout this thesis because of its simplicity, and moreover, it is a protocol which is widely studied and used in the literature, both in game-theoretic and heuristic settings (a non-exhaustive list includes [7, 18, 20–22]).

An important feature that differentiates protocols is their usage and definition of the *deadline* of a negotiation. The *deadline* of a negotiation refers to the time before which an agreement must be reached to achieve an outcome better than the best alternative to a negotiated agreement [23]. Each agent can have its own private deadline, or the deadline can be shared among the agents. The deadline may be specified as a maximum number of rounds [24], or alternatively as a real-time target. Note that when the negotiation happens in real time, the time required to reach an agreement depends on the deliberation time of the agents (i.e., the amount of computation required to evaluate an offer and produce a counter offer).

As in [25, 26], we supplement the alternating-offers protocol with a common global real time line, represented here by $\mathcal{T} = [0, D]$. We stipulate that the deadline has been reached when $t = D$, at which moment both agents receive utility 0.

We represent by $x^t_{A \to B}$ the negotiation outcome proposed by agent A to agent B at time t. A *negotiation thread* or *negotiation trace* (cf. [26, 27]) between two agents A and B at time $t \in \mathcal{T}$ is defined as a finite sequence

$$H^t_{A\leftrightarrow B} := \left(x^{t_1}_{p_1\to p_2}, x^{t_2}_{p_2\to p_3}, x^{t_3}_{p_3\to p_4}, \ldots, x^{t_n}_{p_n\to p_{n+1}}\right),$$

where

1. The offers are ordered over time $\mathcal{T}$: $t_k \leq t_l$ for $k \leq l$.
2. The offers are alternating between the agents: $p_k = p_{k+2} \in \{A, B\}$ for all k.
3. All t_i represent instances of time $\mathcal{T}$, with $t_n \leq t$,
4. The agents exchange complete offers: $x^{t_k}_{p_k\to p_{k+1}} \in \Omega$ for $k \in \{1, \ldots, n\}$.

Additionally, the last element of $H^t_{A\leftrightarrow B}$ may be equal to one of the particles $\{Accept, End\}$. We will say a negotiation thread is *active* if this is not the case.

When agent A receives an offer $x^t_{B\to A}$ from agent B sent at time t, it has to decide at a later time $t' > t$ whether to accept the offer, or to send a counter-offer $x^{t'}_{A\to B}$. Given a negotiation thread $H^t_{A\leftrightarrow B}$ between agents A and B, we can express the action performed by A with a *decision function* [25, 26]. The resulting action is used to extend the current negotiation thread between the two agents. If the agent does not accept the current offer, and the deadline has not been reached, it will prepare a counter-offer by using a negotiation strategy or *tactic* to generate new values for the negotiable issues (see Sect. 2.3).

Various alternative versions of the alternating offers protocol have been used in automated negotiation, extending the default protocol, and imposing additional constraints; for example, in a variant called the *monotonic concession protocol* [15, 28], agents are required to initially disclose information about their preference order associated with each issue and the offers proposed by each agent must be a sequence of concessions, i.e.: each consecutive offer has less utility for the agent than the previous one. Other examples are the three protocols discussed by Fatima et al. [29] that differ in the way the issues are negotiated: simultaneously in bundles, in parallel but independently, and sequentially. The first alternative is shown to lead to the highest quality outcomes. A final example is relevant for our work in Chap. 9 on optimal concession curves, namely a protocol in which only *one* offer can be made. In such a situation, the negotiation can be seen as an instance of the ultimatum game, in which a player proposes a deal that the other player may only accept or refuse [30]. In [31], a similar bargaining model is explored as well; that is, models with one-sided incomplete information and one sided offers. It investigates the role of confrontation in negotiations and uses optimal stopping is to decide whether or not to invoke conflict. The setting of Chap. 9 can also be found in [32], which presents an alternating offer protocol for bilateral bargaining with imperfect information and deadline constraints.

2.2.3 Preference Profiles

Negotiating agents are assumed to have a *preference profile*, which is a *preference order* $\geq$ that ranks the outcomes in the outcome space. Preferences are said to be *ordinal* when they are fully specified by a preference order. Together with the domain they make up the *negotiation scenario*.

An outcome ω' is said to be *weakly preferred* over an outcome ω if $\omega' \geq \omega$. If in addition $\omega \not\geq \omega'$, then ω' is *strictly preferred* over ω, denoted $\omega' > \omega$. An agent is said to be *indifferent* between two outcomes if $\omega' \geq \omega$ and $\omega \geq \omega'$. In that case, we also say that these outcomes are *equally valued* and we write $\omega' \sim \omega$. An *indifference curve* or *iso-curve* is a set of outcomes that are equally valued by an agent. In a *total preference order*, one outcome is always (weakly) preferred over the other outcome for any outcome pair, which means there are no undefined preference relations. Finally, an outcome ω is *Pareto optimal* if there exists no outcome ω' that is preferred by an agent without making another agent worse off [23]. For two players A and B with respective preference orders $\geq_A$ and $\geq_B$, this means that there is no outcome ω' such that:

$$\left(\omega' >_A \omega \wedge \omega' \geq_B \omega\right) \vee \left(\omega' >_B \omega \wedge \omega' \geq_A \omega\right).$$

An outcome that is Pareto optimal is also said to be *Pareto efficient*. When an outcome is not Pareto efficient, there is potential, through re-negotiation, to reach a more preferred outcome for at least one of the agents without reducing the value for the other.

The outcome space can become quite large, which means it is usually not viable to explicitly state an agent's preference for every alternative. For this reason, there are more succinct preference representations for preferences [33, 34].

A well-known and compact way to represent preference orders is the formalism of conditional preference networks (CP-nets) [35–37]. CP-nets are graphical models, in which each node represents an negotiation issue and each edge denotes preferential dependency between issues. If there is an edge from issue i to issue j, the preferences for j depend on the specific value for issue i. To express conditional preferences, each issue is associated with a conditional preference table, which represents a total order of possible values for that issue, given its parents' values.

A preference profile may be specified as a list of ordering relations, but it is more common in the literature to express the agent's preferences by a *utility function*. A utility function assigns a utility value to every possible outcome, yielding a *cardinal* preference structure.

Cardinal preferences are 'richer' than ordinal preferences in the sense that ordinal preferences can only compare between different alternatives, while cardinal preferences allow for expressing the intensity of every preference [33]. Any cardinal preference induces an ordinal preference, as every utility function u defines an order $\omega' \geq \omega$ if and only if $u(\omega') \geq u(\omega)$.

Some learning techniques make additional assumptions about the structure of the utility function [38], the most common in negotiation being that the utility of a multi-issue outcome is calculated by means of a *linear additive function* that evaluates each issue separately [23, 38, 39]. Hence, the contribution of every issue to the utility is linear and does not depend on the values of other issues. The utility $u(\omega)$ of an outcome $\omega = \langle \omega_1, \ldots, \omega_n \rangle \in \Omega$ can be computed as a weighted sum from evaluation functions $e_i(\omega_i)$ as follows:

$$u(\omega) = \sum_{i=1}^{n} w_i \cdot e_i(\omega_i), \tag{2.1}$$

where the w_i are normalized weights (i.e. $\sum w_i = 1$). Linear additive utility functions make explicit that different issues can be of different importance to a negotiating agent and can be used to efficiently calculate the utility of a bid at the cost of expressive power, as they cannot represent interaction effects (or dependencies) between issues [36].

A common alternative is to make use of non-linear utility functions to capture more complex relations between offers at the cost of additional computational complexity. Non-linear negotiation is an emerging area within automated negotiation that considers multiple inter-dependent issues [40, 41]. Typically this leads to larger, richer outcome spaces in comparison to linear additive utility functions. A key factor in non-linear spaces is the ability of a negotiator to make a proper evaluation of a proposal, as the utility calculation of an offer might even prove NP-hard [42]. Examples of this type of work can be found in [43–46].

For non-linear utility functions in particular, a number of preference representations have been formulated to avoid listing the exponentially many alternatives with their utility assessment [33]. The utility of a deal can be expressed as the sum of the utility values of all the *constraints* (i.e., regions in the outcome space) that are satisfied [43, 47]. These constraints may in turn exhibit additional structure, such as being represented by hyper-graphs [48]. One can also decompose the utility function into *subclusters* of individual issues, such that the utility of an agreement is equal to the sum of the sub-utilities of different clusters [46]. This is a special case of a utility structure called *k-additivity*, in which the utility assigned to a deal can be represented as the sum of basic utilities of subsets with cardinality $\leq k$ [49]. For example, for $k = 2$, the utility $u(\omega_1, \omega_2, \omega_3)$ might be expressed as the utility value of the individual issues $u_1(\omega_1) + u_2(\omega_2) + u_3(\omega_3)$ (as in the linear additive case), plus their 2-way interaction effects $u_4(\omega_1, \omega_2) + u_5(\omega_1, \omega_3) + u_6(\omega_2, \omega_3)$. This is in turn closely related to the OR and XOR languages for bidding in auctions [50], in which the utility is specified for a specific set of clusters, together with rules on how to combine them into utility functions on the whole outcome space.

In our setting, both the domain and preferences stay fixed during a single negotiation encounter, but while the domain is common knowledge to the negotiating parties, the preferences of each player is private information. This means that the players do not have access to the utility function of the opponent. In more detail, even the opponent's orderings of the issues are unknown, and the agents are not provided with any prior distribution over the utility functions. However, the players can attempt to learn during the negotiation encounter.

The preference profile of an agent may also specify a *reservation value*. The reservation value is the minimal utility that the agent still deems an acceptable outcome. That is, the reservation value is equal to the utility of the best alternative to no agreement. A bid with a utility lower than the reservation value should not be offered or accepted by any rational agent. In a single-issue domain, the negotiation is often about the *price* P of a good [25, 27, 51, 52]. In that case, agent A and B usually take the roles of buyer B and seller S, and their reservation values are specified by their *reservation prices* RP_B and RP_S. RP_B denotes the highest price a buyer is willing to pay, while RP_S is the lowest price at which a seller is willing to sell.

The negotiator's nearness to a deadline is only one example of *time pressure* [53], which is defined as a negotiator's desire to end the negotiation quickly [54]. Another way to model time pressure is to supplement the negotiation scenario with a *discount factor*. Let d in [0, 1] be the discount factor and let t in [0, 1] be the current normalized time. We compute the discounted utility $u^d(\omega)$ from the undiscounted utility $u(\omega)$ as follows:

$$u^d(\omega) = u(\omega) \cdot d^t. \tag{2.2}$$

If $d = 1$, the utility is not affected by time, and such a scenario is considered to be undiscounted, while if d is very small, there is high pressure on the agents to reach an agreement. Note that discount factors are part of the scenario, are known to both agents and are always *symmetric* (i.e. d always has the same value for both agents).

The reasons for having deadlines and discount factors are both pragmatic and to make the negotiation more interesting from a theoretical perspective. Without a deadline or discount factor, the negotiators have no incentive to accept an offer, and so the negotiation might go on forever. Also, with unlimited time an agent may simply try a large number of proposals to learn the opponent's preferences. In addition, as opposed to having a fixed number of rounds, both the discount factor and deadline are measured in real time. This, in turn, introduces another factor of uncertainty since it is now unclear how many negotiation rounds there will be, and how much time an opponent requires to compute a counter offer. Also, this computational time will typically change depending on the size of the outcome space.

2.2.4 Outcome Spaces

A useful way to visualize the preferences of both players simultaneously is by means of an *outcome space plot* (Fig. 2.2). The axes of the outcome space plot represent the utilities of player A and B, and every possible outcome $\omega \in \Omega$ maps to a point $(u_A(\omega), u_B(\omega))$. The line that connects all of the Pareto optimal agreements is the Pareto frontier.

Note that the visualization of the outcome space together with the Pareto frontier is only possible from an *external* point of view. In particular, the agents themselves are not aware of the opponent utility of bids in the outcome space and do not know the location of the Pareto frontier.

From Fig. 2.2 we can immediately observe certain characteristics of the negotiation scenario. For example, the domain size, whether the bids are spread out over the domain, and the relative occurrence of Pareto optimal outcomes.

One important measure is the *bid distribution*, which is defined as the mean distance to the Pareto frontier. A scenario with a high bid distribution has a high

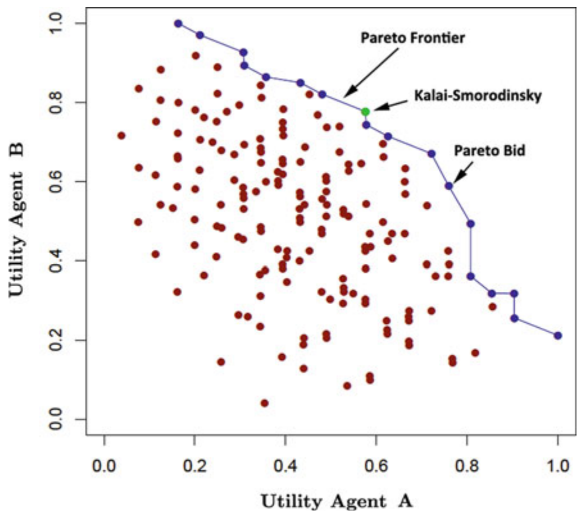

Fig. 2.2 A typical example of an outcome space between agents A and B

percentage of outcomes far from the Pareto frontier. This is defined formally as:

$$\text{distribution}(\Omega) = \sum_{\omega \in \Omega} \frac{\min_{p \in \Omega_P} d(\omega, p)}{|\Omega|}, \tag{2.3}$$

where $\Omega_P \subseteq \Omega$ is the set of Pareto efficient possible outcomes.

There are a number of special outcomes in the outcome space. Of course, the best result would be the outcome $\overline{\omega}$ at which both parties would receive their maximum utility. This would lead to complete satisfaction of both parties, but unfortunately, this is usually not a possible outcome.

There are also a number of definitions for what constitutes a *fair outcome* for both players [23]. The *Nash solution* is defined as the outcome that maximizes the product of the utilities of agents A and B:

$$\omega_{\text{Nash}} = \max_{\omega \in \Omega} u_A(\omega) \cdot u_B(\omega). \tag{2.4}$$

An alternative is the *Kalai-Smorodinsky solution*, which is defined as:

$$\omega_{\text{Kalai}} = \min_{\omega \in \Omega} \left(\frac{u_A(\overline{\omega})}{u_B(\overline{\omega})} - \frac{u_A(\omega)}{u_B(\omega)} \right). \tag{2.5}$$

The opposition of the negotiation scenario is determined by the minimum distance from the Kalai-Smorodinsky solution to the point $\overline{\omega}$.[1] Formally:

[1]There are various ways to define the opposition of a scenario (see [55]), but as in [56], we will employ a definition based on distance measures throughout the thesis. Another popular definition is: $\text{opposition}(\Omega) = \min_{\omega \in \Omega} d(\omega, \overline{\omega})$.

$$\text{opposition}(\Omega) = d(\omega_{\text{Kalai}}, \overline{\omega}) \tag{2.6}$$

where Ω is the set of all possible outcomes, and $d(\omega_1, \omega_2)$ gives the Euclidean distance between two points ω_1, ω_2 in the outcome space, as defined in Eq. (2.7).

$$d(\omega_1, \omega_2) = \sqrt{(u_A(\omega_1) - u_A(\omega_2))^2 + (u_B(\omega_1) - u_B(\omega_2))^2}, \tag{2.7}$$

When a gain for one party can be achieved only at a loss for the other party (i.e., when the preferences are *conflicting*), the negotiation scenario is said to be *competitive*, or to have *strong opposition*. Conversely, in a *cooperative* scenario (or: a scenario with *weak opposition*), both parties achieve either losses or gains simultaneously.

2.3 Negotiating Strategies

A negotiating agent employs a *negotiation strategy* to determine its action in a given negotiation state. Research on general agent negotiators has given rise to a broad variety of negotiation strategies that have already been established both in literature and in implementations, (e.g. [27, 57–61]). The strategies of the agents usually vary from equilibrium strategies in a game theoretical setting to more heuristic approaches. Here we focus in particular on self-interested, boundedly rational agents that are able to conduct bilateral negotiations with incomplete information (following the classification of [16]).

Examples of such general agent negotiators in the literature include, among others: Zeng and Sycara [52], who introduce a generic agent called *Bazaar*; Faratin et al. [58], who propose an agent that is able to make trade-offs in negotiations and is motivated by maximizing the joint utility of the outcome (that is, the agents are utility maximizers that seek Pareto-optimal agreements); Karp et al. [62], who take a game-theoretic view and propose a negotiation strategy based on game-trees; Jonker et al. [60], who propose a a concession oriented strategy called *ABMP*; and Lin et al. [63], who propose an agent negotiator called *QOAgent*.

The ANAC competition that we hosted brought forth an additional 60 advanced negotiation strategies (see Appendix B on ANAC and Appendices C–F for agent descriptions). Notable ANAC agent strategies include: *Agent K* [64, 65], which calculates its target utility based on the average and variance of previous bids and employs a sophisticated acceptance strategy; *IAMHaggler* [66–68], which uses Gaussian process regression technique to predict the opponent's behavior; *CUHK Agent* [69, 70], which adaptively adjusts its acceptance threshold based on domain and opponent analysis; *OMAC Agent* [71–74], which models the opponent using wavelet decomposition and cubic smoothing spline; *The Fawkes*, which combines the best bidding, learning, and accepting strategy components; and finally, *Meta-Agent* [75–77], which, for any given negotiation domain, dynamically selects the most successful ANAC agent to produce an offer.

In Chap. 3, we introduce a component-based architecture for negotiating agents, so we start by describing literature that investigates and evaluates such components. There are two categories of relevant work we highlight here: literature detailing the architecture of a negotiating agent's strategy (Sect. 2.3.1); and work that explores and combines a set of negotiation strategy components to find better strategies (Sect. 2.3.2).

Our component-based architecture consists of three basic components: a *bidding strategy*, which determines which concession should be made in a negotiation state; an *acceptance strategy*, which is used by an agent to determine whether an opponent's offer should be accepted; and optionally an *opponent model*, which can be used both by the *bidding strategy* and *acceptance strategy* to reach a better outcome by exploiting knowledge about the opponent. We provide some background on each of the components in Sects. 2.3.3–2.3.5.

2.3.1 Architecture of Negotiation Strategies

To our knowledge, there is little work in literature describing the generic components of a negotiation strategy architecture, at a similar level of detail as our BOA architecture, which is outlined in Chap. 3. For example, Bartolini et al. [78] and Dumas et al. [79] treat the negotiation strategy as a singular component.

Jonker et al. [60] present an agent architecture for multi attribute negotiation, where each component represents a specific process within the behavior of the agent, e.g.: attribute evaluation, bid utility determination, utility planning, and attribute planning. There are some similarities between the two architectures; for example, the utility planning and attribute planning component correspond to the bidding strategy component in our architecture. In contrast to our work however, Jonker et al. focus on tactics for finding a counter offer and do not discuss acceptance strategies. The fact that our architecture allows this, makes it possible to find better strategies to accept (see Chaps. 4 and 5).

Ashri et al. [80] introduce a general architecture for negotiation agents, discussing components that resemble our architecture; components such as a proposal evaluator and response generator resemble an acceptance condition and bidding strategy respectively. However, the negotiation strategy is described from a BDI-agent perspective (in terms of motivation and mental attitudes).

Hindriks et al. [81] introduce an architecture for negotiation agents in combination with a negotiation system architecture. Parts of the agent architecture correspond to our architecture, but they treat the acceptance strategy and bidding strategy as a singular component, and their focus is primarily on how the agent framework can be integrated into a larger system.

2.3.2 Negotiation Strategy Space Exploration

There are various ways to explore the automated negotiation strategy space by combining a set of negotiation strategies.

Faratin et al. [27] analyze the performance of pure negotiation tactics on single issue domains in a bilateral negotiation setting. The decision function of the pure tactic is then treated as a component around which the full strategy is built. While they discuss how tactics can be linearly combined, the performance of the combined tactics is not analyzed.

Some authors use genetic algorithms to automatically combine certain tactics or strategies. This approach is different to how we combine components using the BOA framework, however they do share certain traits, as they view a strategy consisting of different components and combine them in order to produce a better performing strategy. For example, Matos et al. [82] employ a set of baseline negotiation strategies that are time dependent, resource dependent, and behavior dependent [27], all with varying parameters. The negotiation strategies are encoded as chromosomes and combined linearly, after which they are used by a genetic algorithm to analyze the effectiveness of the strategies. The fitness of an agent is its score in a negotiation competition. This approach analyzes acceptance criteria that only specify a utility interval of acceptable values, and hence do not take time into account. The agents also do not employ explicit opponent modeling.

Eymann [83] also uses genetic algorithms with more complex negotiating strategies, evolving six parameters that influence the bidding strategy. The genetic algorithm uses the current negotiation strategy of the agent and the opponent strategy with the highest average income to create a new strategy, similar to other genetic algorithm approaches (see Beam and Segev [84] for a discussion of genetic algorithms in automated negotiation). The genetic algorithm approach mainly treats the negotiation strategy optimization as a search problem in which the parameters of a small set of strategies are tuned by a genetic algorithm. In Chap. 3, we analyze a more complex space of newly developed negotiation strategies, as our pool of surveyed negotiation strategies consists of strategies introduced in the ANAC competition, as well as the strategies discussed by Faratin et al. (see Sect. 2.3.3). Furthermore, our work combines different *components* instead of complete strategies or strategy parameters and also investigates the importance of particular components (see Chap. 10).

Ros and Sierra obtain promising results in [85] with a negotiation strategy that combines two components: a concession based strategy (either time-based or behavior-based [27]) that decreases a utility threshold to achieve an agreement, and a trade-off strategy [58] that searches for a satisfactory proposal. Our work in this thesis differs with Sierra et al. as we consider a much wider array of agents of which we are able to change the opponent model as well.

Finally, Ilany and Gal [75–77] take the approach of selecting the best strategy from a predefined set of agents, based on the characteristics of a domain. Through machine learning this agent is optimized to choose the best strategy for that particular domain. The difference with our work is that they combine whole strategies, whereas the BOA architecture combines the *components* of strategies. Our contribution is to define and implement an architecture that allows to easily vary all main components of a negotiating agent. Especially in Chap. 10, we study the effects of a much larger group of state of the art negotiation components than has been done before.

Another way to explore the space of negotiation strategies is to classify them according to their behavior. We do so in Chap. 8, in which we present a new classification method for negotiation strategies, based on their pattern of concession making. This chapter is inspired by ideas presented in [86] (of which parts originally appeared in unpublished work by Kersten in 2005). In [86], four dual negotiation orientations are distinguished, depending on the negotiator's own orientation and that of the negotiating partner. Both orientations can be either competitive or cooperative, leading to four different labels: *Competitor*, *Yielder*, *Exploiter*, and *Cooperator*. In Chap. 8, we re-use these labels to name the stance of a negotiator against different kinds of opponents. However in our work, the negotiators are assumed to have different responses to different observed behavior by the other party. Therefore, instead of the negotiator having one particular stance during the negotiation, the position of the negotiators can change in response to the competitiveness of the opponent. For example, a negotiator may be *both* an *Exploiter* (against a *Cooperator*), and a *Yielder* (against a *Competitor*). The negotiator would then be called an *Inverter*, as he takes on the reverse role of his opponent.

In [87], a classification scheme is given for electronic commerce negotiation, including characteristics of the negotiating agents. It is argued that agents can act in a self-interested way, or altruistically, or strike a balance in between. This choice is then seen as a component of the bidding strategy of the agent, which ultimately decides how and when to place offers, or when to withdraw, etc. Although the paper makes this distinction in bidding characteristics, it does not provide a definition or a way to quantify them.

Thomas [88] defines five conflict–handling modes that can be applied to negotiation: *competing*, *collaborating*, *compromising*, *avoiding*, and *accommodating*. Similar to our work in Chap. 8, the classification method uses two underlying dimensions. However, the underlying dimensions are different, namely: *assertiveness* (attempting to satisfy one's own concerns), and *cooperativeness* (attempting to satisfy other's concerns). This classification method is phrased in qualitative, intentional terms of the conflict-handler. Similarly, Zachariassen [89] distinguishes negotiation strategies into two strategy types: *distributive* and *integrative*. This description also focuses on the approach used by the negotiators. Our work has a different focus from both papers, centering around quantitative negotiation characteristics *in response* to agents having either high and low concession rates. Furthermore, we do not classify negotiation strategies in a binary way (either cooperative or non-cooperative), but we employ a continuous spectrum in our approach to classify the full space of negotiation strategies.

2.3.3 *Bidding Strategies*

The *bidding strategy*, also called *negotiation tactic* or *concession strategy*, is usually a complex strategy component. Two types of negotiation tactics are very common: *time-dependent tactics* and *behavior-dependent tactics*. Each tactic uses a *decision function*, which maps the negotiation state to a *target utility*. Next, the agent can search for a bid with a utility close to the target utility and offer this bid to the opponent (Fig. 2.3).

2.3.3.1 Time-Dependent Tactics

Functions which return an offer solely based on time are called *time-dependent tactics*. The standard time-dependent strategy calculates a target utility $u(t)$ at every turn, based on the current time t. Perhaps the most popular time-based decision function can be found in [20, 27], which, depending on the current normalized time $t \in [0, 1]$, makes a bid with utility closest to

$$u(t) = P_{min} + (P_{max} - P_{min}) \cdot (1 - F(t)), \tag{2.8}$$

where

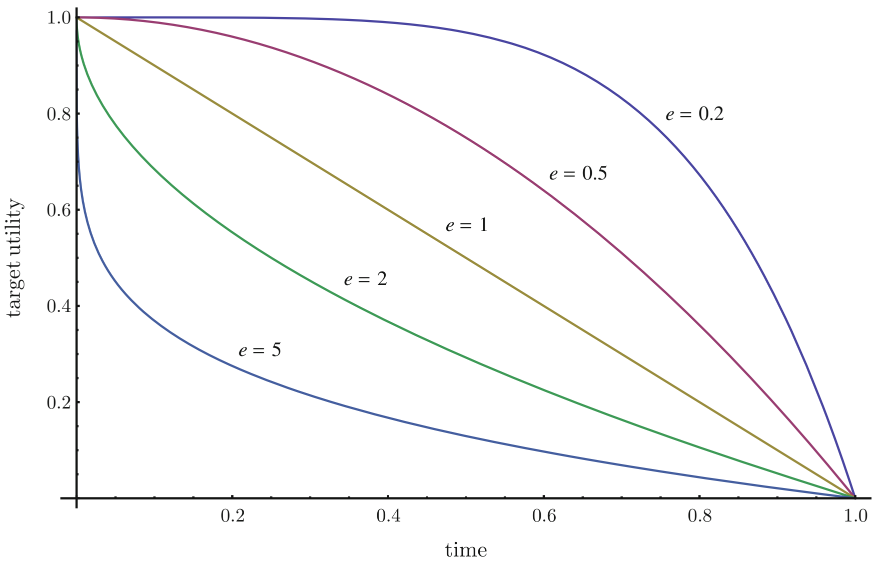

Fig. 2.3 Target utility through time of time-dependent tactics with concession factor $e \in \{0.2, 0.5, 1, 2, 5\}$

$$F(t) = k + (1 - k) \cdot t^{1/e}.$$

The constants $P_{\min}, P_{\max} \in [0, 1]$ control the range of the proposed offers, and $k \in [0, 1]$ determines the value of the first proposal. For $0 < e < 1$, the agent concedes only at the end of the negotiation and is called *Boulware*. If $e \geq 1$, the function concedes quickly to the reservation value, and the agent is then called a *Conceder*.

For $k = 0$ and $e = 1$, we obtain a very simple conceding tactic called *Conceder Linear*. It reduces Eq. (2.8) to

$$u(t) = P_{\min} + (P_{\max} - P_{\min}) \cdot (1 - t),$$

so that the agent linearly reduces its demanded utility from $P_{\max}$ to $P_{\min}$ as time passes.

In many of our experiments in later chapters, we set $k = 0$, and $P_{\max}, P_{\min}$ are respectively set to the maximum and minimum utility that can be obtained in the negotiation scenario. The specification of these strategies given in [20, 27] does not involve any opponent modeling; that is, given the target utility, a random bid is offered with a utility closest to it. Time-dependent tactics accept if and only if the opponent's bid is better than the target utility.

2.3.3.2 Baseline Tactics

The *Hardliner* strategy (also known as *take-it-or-leave-it*, *sit-and-wait* [90] or *Hardball* [91]) can be viewed as an extreme type of time-dependent tactic. This strategy simply makes a bid of maximum utility for itself and never concedes, and is therefore the most competitive strategy that can be implemented.

Random Walker (also known as the *Zero Intelligence* strategy [92]) generates random bids and thus provides the extreme case of a maximally unpredictable opponent. Because of its limited capabilities, it can also serve as a useful baseline strategy when testing the efficacy of other negotiation strategies.

2.3.3.3 Behavior-Dependent Tactics

Faratin et al. introduce a well-known set of *behavior-dependent tactics* or *imitative tactics* in [27]. The most well-known example of a behavior-dependent tactic is the *Tit for Tat* strategy, which tries to reproduce the opponent's behavior of the previous negotiation rounds by reciprocating the opponent's concessions. Thus, *Tit for Tat* is a strategy of cooperation based on reciprocity [93].

Tit for Tat has been applied and found successful in many other games, including the *Iterated Prisoner's Dilemma game* [94]. It is considered to be a very robust strategy, mainly because of the following three features:

1. It is never the first to defect (i.e., it plays nice as long as the the opponent plays nice as well);
2. It can be provoked into retaliation by a defection of the opponent;
3. However, it is forgiving after just one act of retaliation.

In total three tactics are defined: *Relative Tit for Tat*, *Random Absolute Tit for Tat*, and *Averaged Tit for Tat*. The *Relative Tit for Tat* agent mimics the opponent in a percentage-wise fashion by proportionally replicating the opponent's concession that was performed $\delta \geq 1$ steps ago. The decision function of *Relative Tit for Tat* is as follows:

$$x_{a\to b}^{t_{n+1}}[j] = \min\left(\max\left(\frac{x_{b\to a}^{t_{n-2\delta}}[j]}{x_{b\to a}^{t_{n-2\delta+2}}[j]} x_{a\to b}^{t_{n-1}}[j], min_j^a\right), max_j^a\right) \tag{2.9}$$

The formula specifies the value for each issue j for the next bid for the opponent $x_{a\to b}^{t_{n+1}}$ at time step t_{n+1}, and depends on the previous opponent offers $x_{b\to a}^{t_{n-2\delta}}[j]$ and $x_{b\to a}^{t_{n-2\delta+2}}[j]$ in proportion to its own previous offer $x_{b\to a}^{t_{n-1}}[j]$. The min and max functions are used to ensure that the value of each issue stays within the acceptable range. The main weakness of the decision function is that a percentage concession by the opponent on a specific issue is in general unequal in utility compared to the same concession by the agent.

The standard *Tit for Tat* strategies from [27] do not employ any learning methods, but this work has been subsequently extended by the *Nice Tit for Tat agent* [95] and the *Nice Mirroring Strategy* [96]. These strategies achieve more effective results by combining a simple *Tit for Tat* response mechanism with learning techniques to propose offers closer to the Pareto frontier. These approaches can be viewed as simple examples of the ideas we explore in Sect. 3, where we study arbitrary combinations of concession strategies with learning methods.

2.3.4 Acceptance Strategies

All negotiation agent implementations have to deal with the question of when to accept. In many cases, the agent accepts a proposal when the value of the offered contract is higher than the offer it is ready to send out at that moment in time. This is a significant case, in which the bidding strategy effectively dictates the acceptance strategy. Examples include the time-dependent negotiation strategies defined in [85] (e.g. the *Boulware* and *Conceder* tactics). The same principle is used in the equilibrium strategies of [20] and for the *Trade-off agent* [58]. *Agent K* [64] employs a more sophisticated method to decide when to accept. Its *acceptance strategy* (or *acceptance mechanism*) is based on the mean and variance of all received offers. It then tries to determine the best offer it might receive in the future and sets its proposal target accordingly. We refer to the agent descriptions in C.1 and D.1 in the Appendix for more descriptions of acceptance strategies.

We treat acceptance mechanism design in more detail in Chap. 4, where we present a model for accepting offers and where we compare state-of-the-art acceptance conditions of a large set of negotiation strategies. Our negotiation model builds upon the model of [26], where one specific acceptance condition is studied. We take a more general approach in Chaps. 4 and 5, in which the agent utilizes a generic acceptance mechanism where the current time and the entire bidding history is considered.

We only consider the alternating offers protocol in this thesis, but there are multiple other accepting strategies available for other methods of reaching an agreement. In a multi-party setting, the problem of when to accept is more complex, as the outside options become dynamic; however, the presence of a mediator can reduce some of the complexity by taking over the role of finding acceptable agreements, for example through letting the agents vote on whether a proposed contract is acceptable [44]. It may then be sufficient for an agent to simply accept anything above its reservation value. In the same way, when richer protocols are employed (e.g., when communication is possible, for instance in persuasive, or argumentation-based negotiation [8, 97]), the acceptance dilemma may be easier to resolve, as agents have more knowledge about the acceptability of offers. Lastly, in traditional negotiation protocols such as alternating offers, once a contract is settled upon, it is binding. However, a more general approach is to allow decommitment, i.e. backing out of the negotiation after finding a superior option elsewhere, usually at the cost of a penalty [98]. This requires complex acceptance strategies for committing and decommitting to agreements in a concurrent way, which has recently opened up new research in this area [99–101].

2.3.4.1 Optimal Stopping

When we move from real-time negotiation to round-based negotiations, it becomes possible to adopt *optimal* acceptance strategies through backward inductive reasoning; the most well-known solution being that agreement is reached immediately in the first round [19]. In a real-time setting, it is generally unknown when the last offer has been made, and this makes it difficult to find optimal acceptance conditions for this setting.

In Chap. 5 we explore this idea, and we present the first work that deals with the optimal decision on the acceptance of an offer in a negotiation setting of incomplete information. In many settings of *complete* information ([19] is a typical example) the deal is usually formed right away and as such, sequential decisions whether to accept do not come into play. In [52], a sequential decision making framework is also employed, using similar arguments for using it as we do. Furthermore, they also choose actions that maximize the expected payoff using a recursive formula; however, their approach uses Bayesian learning techniques and does not provide solutions specifically aimed at acceptance strategies. The work by Fatima et al. [25] also treats optimal strategies in an incomplete information setting, but it primarily focuses on bidding strategies in the context of unknown deadlines and reservation values, and does not deal with acceptance strategies. Research that comes closest to our work on optimal acceptance strategies is presented in [31], where optimal stopping is employed to decide when a party should reach an agreement in the context of conflict resolution. In contrast to our work, the scope of the paper is limited to simple bargaining games, and deals with one-sided incomplete information only.

We come back to optimal stopping and sequential decision making in Chap. 9 when we formulate optimal concession curves. To the best of our knowledge, that is the first work that makes usage of the optimal stopping rule to *generate* offers in an incomplete information setting and compares it to other concession techniques, where other previous work makes use of optimal stopping theory to formulate acceptance strategies in different settings [102–104]; for instance deciding to accept sequential job offers while trying to maximize the sum of the payments of all accepted jobs [104]. The major difference with optimal acceptance policies and our work in Chap. 9 is that we use the optimal stopping rule for *concessions*, instead of focusing on the optimal time to accept. Our work in Chap. 9 is defined more as the *complimentary* version of our approach in Chap. 5, in the sense that our formulation of optimal bidding rules happen to resemble optimal acceptance rules. Another key point is that we do not assume that the players' strategies are fixed, which allows us to formulate optimal bidding strategies against certain types of accepting strategies.

2.3.5 Opponent Models

An opponent model is an abstracted description of a player (and/or of a player's behavior) during the game [105]. There are many different types of opponent models; for instance, a model can describe the opponent's preferences, strategy, weaknesses, knowledge, and so on. We present here a short background on learning techniques and evaluation techniques in negotiation for our setting; for a more detailed exposition we refer to our survey on this topic [106].

In negotiation, opponent modeling often revolves around three questions:

- **Preference estimation**. What does the opponent want?
- **Strategy prediction**. What will the opponent do, and when?
- **Opponent classification**. What kind of player is the opponent, and how should we act accordingly?

The above questions are often highly related. For example, some form of preference estimation is needed in order to adequately interpret the opponent's actions. Then, knowing how the opponent acted according to its own utility, we can deduce its strategy, which in turn can help predict what the agent will do in the future. We will mainly focus on preference modeling in this thesis, although our architecture can accommodate for the other types of opponent models as well (see Chaps. 6 and 7).

Constructing an opponent model may alternatively be viewed as a *classification problem* where the *type* of the opponent needs to be determined from a range of possibilities [107]; one example being the work by Lin et al. [63]. Here the *type* of an opponent refers to all opponent attributes that may be modeled to gain an advantage in the game.

Opponent modeling can be performed online or offline, depending on the availability of historical data. Offline models are created before the negotiation starts, using previously obtained data from earlier negotiations. Online models are constructed from knowledge that is collected during a single negotiation session, which is the focus of this thesis. A major challenge in online opponent modeling is that the model needs to be constructed from a limited amount of exchanged bids, and real-time deadlines may pose the additional challenge of having to construct the model as fast as possible. Even though there are large differences between the models, a common set of high level motivations behind their construction can be identified. There are the following motivations for why opponent models are used in automated negotiation:

1. **Augment behavior-based tactics** [95, 96, 108, 109] An opponent model can assist in improving the performance of behavior-based tactics, as the opponent's concessions can then be estimated and reciprocated more accurately. Based on move classification, behavior-based strategies such as the *Tit for Tat* strategy (see Sect. 2.3.3.3) can be applied. In addition, in a negotiation where the opponent's preferences are private, an agent's concession might accidentally result in a decrease in utility for the opponent as well. Such an offer is called *unfortunate* [108], and can be avoided by better estimating the opponent's preferences.
2. **Avoid non-agreement** [28, 110–120] In most negotiations, reaching an agreement is preferred over not reaching a deal. The opponent's previous moves can be analyzed to estimate the minimal concessions required to ensure acceptance.
3. **Find a counter-strategy** [68, 69, 72, 73, 110, 111, 121–133] The opponent can be exploited in multiple ways with the assistance of an opponent model. One way is to estimate the opponent's reservation value in an attempt to obtain the minimal negotiation outcome the opponent will settle for. Alternatively, an estimate of the opponent's deadline can be used to elicit concessions from the opponent by stalling the negotiation, provided of course that the agent has a later deadline. Theoretical results are available that specify which counter-strategy to use depending on the information known about the opponent [25, 82, 134].
4. **Maximize social welfare** [46, 51, 52, 63, 112–115, 118–120, 135–141] In a cooperative environment, agents aim for a fair result. An agent can use an estimate of the opponent's preference profile to maximize the chances of a good outcome for both.
5. **Propose Pareto optimal bids** [28, 46, 67, 95, 96, 109, 118, 135, 137, 140, 142–152] Pareto optimality of an offer ensures the offer cannot be improved for both players at the same time. When an agent considers multiple similarly preferred offers to send out to the opponent, offering a Pareto optimal bid can lead to an earlier and mutually beneficial agreement.
6. **Reduce negotiation costs** [46, 51, 52, 109, 110, 113, 115, 119, 120, 135–137, 139–141, 148, 153–155] In general it costs time and resources to negotiate, and using an estimate of the opponent's preference profile or negotiation strategy

can aid in reducing these costs. An agent may even decide that the estimated negotiation costs are too high to warrant a potential agreement, and prematurely end the negotiation.

We found that existing work on opponent models can fulfill any of the goals above by learning a combination of *six* opponent attributes, which we have listed in Table 2.1. The notion of an opponent model as a component of a negotiation strategy has been discussed by many of these authors. However, to our knowledge, there is limited work in which the performance of different types of opponent models is compared as we do in Chaps. 6 and 7. One example is the work by Papaioannou et al. [116], who evaluate a set of opponent strategy prediction techniques in terms of resulting performance gain.

2.4 Evaluation Methodologies

We now introduce the methodologies we use in subsequent chapters to evaluate negotiation strategies. The first evaluation method is analytical software to analyze the performance and dynamics of agents, and the outcome of the negotiation (Sect. 2.4.1). The second is a method to benchmark and objectively evaluate negotiation agents in a competitive setting (Sect. 2.4.2). Together, they provide an environment to apply a range of *performance measures* (Sect. 2.4.3) to measure the performance of a negotiation strategy. Lastly, we discuss measures for learning methods, including *accuracy measures* (Sect. 2.4.4).

2.4.1 Environments for Evaluating Negotiating Agents

As we have built a generic environment for designing and evaluating agent negotiators called Genius [165] (see Appendix A), we briefly review related work that is explicitly aimed at the evaluation of various agent negotiators. Most of the work reported herein concerns the evaluation of various *strategies* for negotiation used by such agents. Although some results were obtained by game-theoretic analysis (e.g. [7, 15]), most results were obtained by means of *simulation* (e.g. [166–168]). Devaux and Paraschiv [166] present work that compares agents negotiating in internet agent-based markets. In particular, they compare a strategy of their own agent with behavioral based strategies taken from the literature [27]. The simulations are performed in an abstract domain where agents need to negotiate the price of a product. Similarly, Henderson et al. [168] present results of the performance of various negotiation strategies in a simulated car hire scenario. Finally, Matos et al. [82] conducted experiments to determine the most successful strategies using an evolutionary approach in an abstract domain called the *service-oriented domain*. Even though several of the approaches use an abstract domain with a range of parameters

Table 2.1 An overview of learning techniques and methods that help to learn six different opponent attributes

Opponent attributes	Procedure	Learning techniques
Reservation value	Bidding strategy estimation	Bayesian learning [51, 52, 110, 119, 120, 130, 133, 141]
		Non-linear regression [111, 121, 125, 133]
Deadline	Bidding strategy estimation	Bayesian learning [110, 119, 133]
		Non-linear regression [110, 119, 125, 133]
Issue preference order	Measuring similarity between offers	Bayesian learning [28]
		Kernel density estimation [136, 156]
		Heuristics [143, 147]
	Knowledge of bidding strategy	Simplified genetic algorithm [112]
Outcome preference order	Classification	Bayesian learning [63, 67, 95, 109, 138, 146, 150, 157, 158]
	Data mining aggregate preferences	Random variable estimation [113, 140]
		Graph theory [46, 148]
		Bayesian network [149]
	Logical reasoning and heuristics	Heuristics [69, 144, 145, 151–155, 159]
Bidding strategy	Regression analysis	Non-linear regression [111, 116, 121, 123, 125, 133, 160]
		Polynomial interpolation [116]
		Genetic algorithms [117]
		Bayesian networks [127]
	Time series forecasting	Derivatives [122, 139]
		Signal processing [68, 72, 73, 129, 132, 161]
		Neural networks [114, 116–118, 126, 162–164]
		Markov chains [128]
Acceptance strategy	Interpolation of acceptance probability	Polynomial interpolation [131]
		Kernel density estimation [115]
		Bayesian learning [137]
		Neural networks [124]

that may be varied, we argue that the focus on a single domain in most simulations is restrictive. A similar argument to this end has been put forward in [56]. The analysis of agent negotiators in multiple domains may significantly improve the performance of such agents.

Manistersky et al. [169] discuss how people who design agent negotiators change their design over time. They study how students changed their design of a trading agent that negotiates in an open environment. After initial design of their agents, human designers obtained additional information about the performance of their agents by receiving logs of negotiations between their agents and agents designed by others. These logs provided the means to analyze the negotiation behavior, and an opportunity to improve the performance of the agents. The GENIUS environment discussed in Appendix A provides a tool that supports such analysis, subsequent improvement of the design, and structures the enhancement process.

Part of GENIUS' functionality has been described in [170, 171], and our work [165] outlined in Appendix A is a natural extension of this research. Since then, we have extended GENIUS with all ANAC resources and new functionality described in Appendix B (e.g., negotiation strategies, protocols, scenarios, discount factors, reservation values), the BOA architecture and agent components from Chap. 3, the acceptance strategies from Chaps. 4 and 5, and the performance and accuracy measures described in Chaps. 6 and 7.

With regard to systems that facilitate the actual design of agents or agent strategies in negotiations, few systems are close to GENIUS. Most of the systems that may be related to its main focus are negotiation support systems (e.g., the Interactive Computer-Assisted Negotiation Support system (ICANS) presented in [172], the InterNeg Support Program for Intercultural REsearch (INSPIRE)), however, GENIUS advances the state-of-the-art by also providing evaluation mechanisms that allow a quick and simple evaluation of strategies and the facilitation of automated negotiator's design. INSPIRE, by Kersten and Noronha [173], is a Web-based negotiation support system with the primary goal of facilitating negotiation research in an international setting. The system enables negotiation between two people, collects data about negotiations and has some basic functionality for the analysis of the agreements, such as calculation of the utility of an agreement and exchanged offers. However, unlike GENIUS, it does not allow integration of an automated negotiating agent and thus does not include repositories of agents as we propose. Perhaps Neg-o-Net [174] is more similar to GENIUS than all the other support systems. The Neg-o-Net model is a generic agent-based computational simulation model for capturing multi-agent negotiations concerning resource and environmental management decisions. The Neg-o-Net model includes both a negotiation algorithm and some agent models. An agent's preferences are modeled using digraphs (scripts). Nodes represent states of the agent that can be achieved by performing actions (arcs). Each state is evaluated using utility functions. The user can modify the agent's script to model his/her preferences w.r.t. states and actions. While Neg-o-Net is similar to GENIUS, there are at least two important differences. First, they currently do not support the incorporation of human negotiators, but only automated ones. Second, they do not provide any evaluation mechanism of the strategies as GENIUS provides.

A recent development worth noting is the *Negowiki* project [17, 175], which aims to unify current approaches in negotiation research by creating a collection of standardized negotiation scenarios. *Negowiki* is an online framework where researchers can share negotiation scenarios and results. As in GENIUS, analysis of the results is provided, so that researchers can compute a set of metrics over the results of the negotiation (e.g. Pareto optimality, fairness; we elaborate more on this in Sect. 2.4.3). All scenarios offered by *Negowiki* are also available for download in GENIUS format.

2.4.2 Negotiating Agent Competitions

A *competition* can act as a useful and open benchmarking tool to evaluate and compare negotiation agents, as evidenced by successful competitions to advance the state-of-the-art in artificial intelligence such as the *Computer Poker Competition* [176], the Iterated Prisoner's Dilemma game [94] and the *Trading Agent Competition* [177]. Following in their footsteps, we organized four annual instances of the International Automated Negotiating Agents Competition (ANAC).

We elaborate on the goals and results of the competition in Appendix B. Here, we provide a short description of related competitions and outline the differences with ANAC.

2.4.2.1 The Trading Agent Competition

Four games of the *Trading Agent Competition* (TAC) relate to automated negotiating agents [177–181], and some elements of TAC have similar challenges as posed by ANAC:

TAC SCM *TAC Supply Chain Management* was designed to simulate a dynamic supply chain environment. Agents have to compete to secure customer orders and components required for production. In order to do so, the agents have to plan and coordinate their activities across the supply chain. Participants face the complexities of supply chains, which admits a variety of bidding and negotiation strategies.

TAC Ad Auctions In the *TAC Ad Auctions*, game entrants design and implement bidding strategies for advertisers in a simulated sponsoring environment. The agents have to bid against each other to get an ad placement that is related to certain keyword combinations in a web search tool. The advertiser strategies have to decide which keywords to bid on, and what prices to offer. Therefore, the strategies have to optimize their data analysis and bidding tactics to maximize their profit.

TAC Market Design *TAC Market Design* or *The CAT Competition* is a reverse of the normal TAC game: as an entrant you define the rules for matching buyers and sellers, while the trading agents are created by the organizers of the competition. Entrants have to compete against each other to build a robust market mechanism that attracts buyers and sellers.

Power TAC Having started in 2011, *Power TAC* is a fairly recent addition to the TAC games. It is built around a competitive market simulation platform with the goal to direct policy making and to develop and validate intelligent agent technology for trading. It models a electrical energy market, where competing business entities offer energy services to customers.

The challenges posed by TAC are similar as in ANAC, especially the games of *TAC Ad Auctions* and *Power TAC*. The games of TAC can get very complex and the domains of the games are specifically chosen to model a certain scenario of a trading agent problem. Contrastingly, the entrants of ANAC have to consider very generic negotiation domains when they design their agents. On the one hand, this makes ANAC very accessible, as there are no domain-dependent details the participants have to know about. On the other hand, it is very difficult to develop an agent that negotiates well under such a wide variety of circumstances, especially with the unique challenges ANAC poses, which include one-shot bilateral negotiations with a real timeline, combined with incomplete information of the opponent's preferences.

2.4.2.2 The Agent Reputation Trust Competition

The *Agent Reputation Trust Competition* (ART) [182, 183] is also a negotiating agent competition with a testbed that allows the comparison of different strategies. The ART competition simulates a business environment for software agents that use the reputation concept to buy advices about paintings. Each agent in the game is a service provider responsible for selling its opinions when requested. The agent can exchange information with other agents to improve the quality of their appraisals. The challenge is to perceive when an agent can be trusted and to establish a trustworthy reputation. Compared to ANAC, the focus of ART is more on trust: the goal is to perceive which agents can be trusted in a negotiation process and what reputation should be attributed to each agent.

2.4.3 Evaluating Performance of Negotiation Strategies

The ultimate aim of a negotiation strategy is to increase overall performance of the negotiation, which is why *performance measures* are used to evaluate a negotiator's success. Performance measures evaluate the quality of the outcome, usually measured in utility gain, or distance of the agreement to the Pareto frontier. With this method, the success of an opponent model is expressed in terms of the negotiation *result* (as opposed to the whole negotiation *process*; for this we refer to Sect. 2.4.4) The paragraphs below provide an overview of the performance measures in related work.

Average utility. Average utility is by far the most popular performance measure and is used by many authors (e.g., [28, 63, 69, 72, 73, 96, 108–111, 115, 117, 119, 120, 122–126, 128, 129, 131–133, 135, 136, 138, 140–142, 146, 147, 150, 151, 156]). A common application is to consider the average utility of an agent with and without opponent model against a group of opponents on several domains (see for example [28, 69, 73]). Note that the average utility of an agent directly depends on the negotiation setting (as we will see in following chapters), which therefore should be chosen with care.

Distance to a fair outcome. Other authors are concerned with achieving a fair outcome [96, 108, 135, 139], which is especially important if there will be future negotiations between the parties.

Distance to a fair outcome is then calculated as the average Euclidean distance to a fair solution (as defined in Sect. 2.2.4), such as the distance to Nash solution [96, 108, 135, 139] or distance to Kalai-Smorodinsky [96, 108, 135]. As with the average utility measure, the negotiation setting strongly influences the result [184, 185].

Distance to Pareto frontier. An opponent model of the opponent's preferences aids in identifying Pareto optimal bids. For this type of model—assuming it is applied by a bidding strategy that takes the opponent's utility into account—the distance to the nearest Pareto optimal bid directly correlates with the model's quality (see for example [28, 109, 118, 135, 137]). Minimizing this distance to the Pareto-optimal frontier improves fairness and the probability of acceptance.

Joint utility. An alternative method to measure the fairness of an outcome is to calculate the joint utility [51, 52, 63, 112–115, 118–120, 136–139, 141]. The majority of the authors simply use the sum of the utility of the final outcome for the agents (see for example [63, 138]). An alternative used by several authors [51, 52, 114, 141] is to consider the normalized joint utility:

$$u_{\text{joint}} = \frac{(P - RP_S)(RP_B - P)}{(RP_B - RP_S)^2}. \tag{2.10}$$

In this equation, P is the agreed upon price, and RP_B and RP_S are the reservation prices of the buyer and seller respectively. Note that this definition is only applicable to single-issue negotiations. An alternative measure for multi-issue negotiations used by Jazayeriy et al. [112] is the geometric mean:

$$u_{\text{joint}} = \sqrt{u_A \cdot u_B}, \tag{2.11}$$

where u_A and u_B are the utilities achieved by the agents. An attractive property of this metric is that when the utilities are highly unbalanced, this formula better reflects unfairness than by simply calculating the sum of the utilities.

Percentage of agreements. An opponent model may lead to better bids being offered to the opponent, possibly avoiding non-agreement. In situations where an agreement is always better than no agreement, the percentage of agreements is a direct measure of success (see for example [28, 46, 110–121, 127, 133, 140, 144, 148, 155]). An important disadvantage is that the acceptance ratio does not capture the quality of the agreement. Agrawal and Chari, Buffett et al., and Mudgal and Vassileva use a related measure in which they calculate how often one agent outperforms the other with regard to the final outcome [121, 127, 144]. A disadvantage of this method is that an agent might outperform other agents, but still reach a bad outcome. An alternative metric is applied by Robu and Poutré [46, 148], which calculates how often an outcome is reached that maximizes social welfare.

Time of agreement. Various authors measure the duration of the negotiation or the communication load (e.g. [46, 51, 52, 109, 110, 113–115, 119, 120, 135–137, 139–141, 148, 153–155]), because in practical settings there is often a non-negligible cost associated with both. Opponent models can lead to earlier agreements, and thereby reduce costs. An important disadvantage of this metric is that while an opponent model may lead to an earlier agreement, the quality of the outcome for the agent might be lower.

Trajectory analysis. The quality of bidding strategies can be measured by analyzing the percentage and relative frequency of certain types of moves [96, 108, 109]. For example, unfortunate moves are offers that decrease the utility for both agents at the same time. Theoretically, a perfect opponent model of the opponent's preferences would allow an agent to prevent any such unfortunate moves. A disadvantage of this method is that it highly depends on the concession strategy that is used in combination with the opponent model.

2.4.4 Evaluating Learning Methods

The performance measures discussed in Sect. 2.3.5 are benchmarks for an entire negotiation strategy, but they are also often used to test the efficacy of one specific component, the most prevalent being the learning component. The simplest approach is to compare a novel learning technique with a set of baseline strategies. In [146] for example, the performance of the opponent model is estimated by embedding it in a strategy and comparing the average utility against two baseline strategies. The modeling technique discussed by [63] introduces a model for a similar protocol, but in this case the baseline is set by humans. Zeng and Sycara [52] measure performance in terms of social welfare, but focus on single-issue negotiations in which they compare the performance of three settings: both learn, neither learn, and only the buyer learns. Finally, [158] evaluates the accuracy of a model against simple baseline strategies in terms of the likelihood that the correct class is estimated to which the opponent's preference profile belongs.

The performance of an opponent model can also be tested against other models or against a theoretical lower or upper bound, as we do in Chap. 6. For example, Coehoorn and Jennings [136] evaluate the performance of their opponent model using a standard bidding strategy that can be used both with and without a model. The performance of the strategy is evaluated in three settings: without knowledge, with perfect knowledge, and when using an offline opponent model. This work is similar to our work in Chap. 6, however, it differs in the fact that we focus on *online* opponent modeling, and our setting is especially challenging as it involves the time/exploration trade-off. Another example is the work by [113], which introduces two opponent models for e-recommendation in a multi-object negotiation.Finally, [56] defines two accuracy measures and uses these measures to analyze the accuracy of two opponent models. The main differences are that in Chap. 6, we focus on the more general type of multi-issue negotiations, we focus on a larger set of performance measures, and pay more attention to the factors that influence the performance of the model. Furthermore, as far as we know, our work is the first to compare and analyze such a large set of state-of-the art models of the opponent's preference profile.

2.4.4.1 Accuracy Measures

As performance measures are only indirect measures of the negotiating agent's quality, other measures, such as *accuracy measures* can also be included for the purpose of benchmarking learning techniques. Accuracy measures are direct measures of learning quality, as they quantify the difference between the *estimate* and the *estimated*; i.e., they determine the quality of a model by quantifying how well the opponent model *represents* the real preferences of the opponent. An example is the correlation between the estimated and the real outcome space, or the percentage of correctly inferred Pareto optimal outcomes. We describe here the accuracy measures for preference modeling methods, as we come back to them in Chap. 7, where we compare the accuracy of various preference modeling techniques, using established accuracy measures.

For example, Carbonneau et al. [162] calculate the Pearson correlation between the real and estimated utility of the opponent's next bid. Hindriks and Tykhonov [56] extend this approach by measuring the Pearson correlation of the whole outcome space and discuss analogous definitions for the ranking distance. Our method in Chap. 7 incorporates both measures. An alternative approach is to measure the distance between elements of two preference profiles. For example Jazayeriy et al. [112] introduce such measures for the learning error of issue weights. We have incorporated these measures in our method, and we also apply the same measures to quantify the similarity between two full bid spaces.

Finally, there exist accuracy measures tailored to specific learning methods. Buffett and Spencer [157] for example, define a metric for opponent models that use Bayesian learning. The measure is defined as the average likelihood that the correct hypothesis is chosen from the set of candidate hypotheses. Since we employ models in Chaps. 6 and 7 that are based on a wide range of learning techniques, we do not incorporate measures specific to a particular learning method.

We also quantify the *relationship between accuracy and performance* in Chap. 7. In related work by Coehoorn and Jennings [136], a model is introduced that estimates the opponent's issue weights and the influence of small prediction errors on performance is investigated. The method in this thesis takes this a step further, as we analyze the relation between an exhaustive set of accuracy measures—including accuracy of the issue weights—and performance.

Similarity between issue weights We can measure the accuracy of models that estimate the issue weights of the opponent's preference profile in several ways [56, 112, 142]. All of them use a distance metric between the issue weights $w = (w_1, \ldots, w_n)$ of the real opponent preferences u_{op} and the issue weights $w' = (w'_1, \ldots, w'_n)$ of the estimated preferences u'_{op}. One way to do so is to measure the distance between the issue weight vectors [112]:

$$d_{\text{Euclidean}}(w, w') = \sqrt{\sum_{i=1}^{n}(w_i - w'_i)^2}. \tag{2.12}$$

Of course, this measure can be used for scalars as well. When modeling the opponent's deadline (or reservation value) $x \in \mathbb{R}$ with an estimate x', Eq. (2.12) simplifies to

$$d_{\text{Euclidean}}(x, x') = |x - x'|. \tag{2.13}$$

Another way is to check whether the issue weights are *ranked* correctly [56] by evaluating all possible pairs of issues $i_1, \ldots, i_n$:

$$d_{\text{rank}}(w, w') = \frac{1}{n^2}\sum_{j=1}^{n}\sum_{k=1}^{n} c(i_k, i_j), \tag{2.14}$$

where $c(i_k, i_j)$ is the conflict indicator function, which is equal to one when the ranking of the weights of issues i_k and i_j differs between the two profiles, and zero otherwise. An alternative is to measure the *correlation* between the vectors [56]:

$$d_{\text{Pearson}}(w, w') = \frac{\sum_{i=1}^{n}(w_i - \overline{w})(w'_i - \overline{w'})}{\sqrt{\sum_{i=1}^{n}(w_i - \overline{w})^2 \sum_{i=1}^{n}(w'_i - \overline{w'})^2}}. \tag{2.15}$$

Note that this expression may be undefined, for example when all weights are equal.

Similarity between preference profiles When opponent models estimate the opponent's preferences fully (e.g. [56, 142, 148, 157–159]), the quality of these models depends on the similarity between the real u_{op} and estimated opponent's preference profile u'_{op} for all bids in the outcome space Ω. One approach is to calculate the *average distance* between all outcomes in Ω [142]:

$$d_{\text{abs}}(u_{op}, u'_{op}) = \frac{1}{|\Omega|}\sum_{\omega \in \Omega} |\omega - \omega'|. \tag{2.16}$$

However, in practice, the correct ranking of the bids can be sufficient already. An alternative is therefore to use the *ranking distance of bids* measure that compares all preference orderings in a pairwise fashion [56]:

$$d_{\text{rank}}(u_{op}, u'_{op}) = \frac{1}{|\Omega|^2} \sum_{\omega \in \Omega, \omega' \in \Omega} c_{\prec_u, \prec_{u'}}(\omega, \omega'), \tag{2.17}$$

where $c_{\prec_u, \prec_{u'}}$ is the conflict indicator function, which is equal to one when the ranking of the outcomes ω and ω' differs between the two profiles, and zero otherwise. Identically, Buffett et al. count the amount of correctly estimated preference relations [157, 158]. A disadvantage of these approaches is their scalability because all possible outcome pairs need to be compared. This problem can be overcome by using a Monte Carlo simulation; however, a more efficient solution can be to use the *Pearson correlation of bids* [56], which is defined as follows:

$$d_{\text{Pearson}}(u_{op}, u'_{op}) = \frac{\sum_{\omega \in \Omega} (u_{op}(\omega) - \overline{u_{op}})(u'_{op}(\omega) - \overline{u'_{op}})}{\sqrt{\sum_{\omega \in \Omega} (u_{op}(\omega) - \overline{u_{op}})^2 \sum_{\omega \in \Omega} (u'_{op}(\omega) - \overline{u'_{op}})^2}} \tag{2.18}$$

A downside of this measure, although unlikely to occur in practice, is that it is not defined for all inputs, for example when all bids are estimated to have the same utility.

This chapter is based on the following publications: [95, 106, 165, 184–188]

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. A tit for tat negotiation strategy for real-time bilateral negotiations. In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 229–233. Springer Berlin Heidelberg, 2013

Tim Baarslag, Mark J.C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, pages 1–50, 2015

Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen V. Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014

Tim Baarslag, Katsuhide Fujita, Enrico H. Gerding, Koen V. Hindriks, Takayuki Ito, Nicholas R. Jennings, Catholijn M. Jonker, Sarit Kraus, Raz Lin, Valentin Robu, and Colin R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence*, 198:73–103, May 2013

Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. The first automated negotiating agents competition (ANAC 2010). In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations*, volume 383 of *Studies in Computational Intelligence*, pages 113–135, Berlin, Heidelberg, 2012. Springer-Verlag

Tim Baarslag, Koen V. Hindriks, Mark J.C. Hendrikx, Alex S.Y. Dirkzwager, and Catholijn M. Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In Ivan Marsa-

Maestre, Miguel A. Lopez-Carmona, Takayuki Ito, Minjie Zhang, Quan Bai, and Katsuhide Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 61–83. Springer, Japan, 2014

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. Towards a quantitative concession-based classification method of negotiation strategies. In David Kinny, Jane Yung-jen Hsu, Guido Governatori, and Aditya K. Ghose, editors, *Agents in Principle, Agents in Practice*, volume 7047 of *Lecture Notes in Computer Science*, pages 143–158, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg

Tim Baarslag and Koen V. Hindriks. Accepting optimally in automated negotiation with incomplete information. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 715–722, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems

References

1. Kraus S, Wilkenfeld J, Harris MA, Blake E (1992) The hostage crisis simulation. Simul Gaming 23(4):398–416
2. Sandholm T, Lesser VR (1995) Issues in automated negotiation and electronic commerce: Extending the contract net framework. In: Proceedings of the first international conference on multi-agent systems (ICMAS). San Francisco, pp 328–335
3. Shehory O, Kraus S (1998) Methods for task allocation via agent coalition formation. Artif Intell 101(1–2):165–200
4. Smith RG (1980) The contract net protocol: High-level communication and control in adistributed problem solver. IEEE Trans Comput 29(12):1104–1113
5. Deutsch M, Coleman PT, Marcus EC (2000) The handbook of conflict resolution: theory and practice, 1st edn. Jossey-Bass, San Francisco
6. Zlotkin G, Rosenschein JS (1991) Cooperation and conflict resolution via negotiation among autonomous agents in noncooperative domains. IEEE Trans Syst Man Cybern 21(6):1317–1324
7. Kraus S (2001) Strategic negotiation in multiagent environments. MIT Press, Cambridge
8. Sycara KP (1993) Machine learning for intelligent support of conflict resolution. Decis Support Syst 10(2):121–136
9. Jennings NR, Faratin P, Lomuscio AR, Parsons S, Wooldridge MJ, Sierra C (2001) Automated negotiation: Prospects, methods and challenges. Group Decis Negot 10(2):199–215
10. Sycara KP (1985) Arguments of persuasion in labour mediation. Proceedings of the 9th international joint conference on artificial intelligence, vol 1. Morgan Kaufmann Publishers Inc, San Francisco, pp 294–296
11. Sycara KP (1988) Resolving goal conflicts via negotiation. In: Proceedings of the 7th national conference on artificial intelligence. St. Paul, MN, pp 245–250, 21–26 August 1988
12. Robinson WN (1990) Negotiation behavior during requirement specification. In: Proceedings of the 12th international conference on software engineering, pp 268–276
13. Rosenschein JS (1996) Rational interaction: cooperation among intelligent agents. Ph.D. thesis, Stanford University, Stanford, 1986
14. Klein M, Lu SC-Y (1989) Conflict resolution in cooperative design. Artif Intell Eng 4(4):168–180
15. Rosenschein JS, Zlotkin G (1994) Rules of encounter: designing conventions for automated negotiation among computers. MIT Press, Cambridge
16. Lomuscio AR, Wooldridge MJ, Jennings NR (2003) A classification scheme for negotiation in electronic commerce. Group Decis Negot 12(1):31–56

17. Marsa-Maestre I, Klein M, Jonker CM, Reyhan A (2013) From problems to protocols: towards a negotiation handbook. Decis Support Syst
18. Osborne MJ, Rubinstein A (1994) A course in game theory, 1st edn. MIT Press, Cambridge
19. Rubinstein A (1982) Perfect equilibrium in a bargaining model. Econometrica 50(1):97–109
20. Fatima SS, Wooldridge MJ, Jennings NR (2002) Multi-issue negotiation under time constraints. AAMAS'02: Proceedings of the first international joint conference on autonomous agents and multiagent systems. ACM, New York, pp 143–150
21. Kraus S, Wilkenfeld J, Zlotkin G (1995) Multiagent negotiation under time constraints. Artif Intell 75(2):297–345
22. Osborne MJ, Rubinstein A (1990) Bargaining and markets (Economic theory, econometrics, and mathematical economics). Academic Press, London
23. Raiffa H (1982) The art and science of negotiation: How to resolve conflicts and get the best out of bargaining. Harvard University Press, Cambridge
24. Sofer I, Sarne D, Hassidim A (2012) Negotiation in exploration-based environment. In: Proceedings of the twenty-sixth AAAI conference on artificial intelligence
25. Fatima SS, Wooldridge MJ, Jennings NR (2002) Optimal negotiation strategies for agents with incomplete information. Revised papers from the 8th international workshop on intelligent agents VIII, ATAL'01. Springer, London, pp 377–392
26. Sierra C, Faratin P, Jennings NR (1997) A service-oriented negotiation model between autonomous agents. In: Boman M, van de Velde W (eds) Proceedings of the 8th European workshop on modelling autonomous agents in multi-agent world, MAAMAW-97. Lecture notes in artificial intelligence, vol 1237. Springer, Heidelberg, pp 17–35
27. Faratin P, Sierra C, Jennings NR (1998) Negotiation decision functions for autonomous agents. Robot Auton Syst 24(3–4):159–182
28. Niemann C, Lang F (2009) Assess your opponent: A bayesian process for preference observation in multi-attribute negotiations. In: Ito T, Zhang M, Robu V, Fatima S, Matsuo T (eds) Advances in agent-based complex automated negotiations, Studies in computational intelligence, vol 233. Springer, Berlin, pp 119–137
29. Fatima SS, Wooldridge MJ, Jennings NR (2006) Multi-issue negotiation with deadlines. J Artif Intell Res 27:381–417
30. Slembeck T (1999) Reputations and fairness in bargaining-experimental evidence from arepeated ultimatum game with fixed opponents. Experimental, EconWPA
31. Santiago Sánchez-Pagés. The use of conflict as a bargaining tool against unsophisticated opponents. ESE Discussion Papers 99, Edinburgh School of Economics, University of Edinburgh, 2004
32. An B, Gatti N, Lesser VR (2013) Bilateral bargaining with one-sided uncertain reserve prices. Auton Agents Multi-Agent Syst 26:420–455
33. Chevaleyre Y, Dunne PE, Endriss U, Lang J, Lemaître M, Maudet N, Padget J, Phelps S, Rodríguez-Aguilar JA, Sousa P (2006) Issues in multiagent resource allocation. Informatica 30:3–31
34. Domshlak C, Hüllermeier E, Kaci S, Prade H (2011) Preferences in AI: an overview. Artif Intell 175(7–8):1037–1052
35. Aydoğan R, Baarslag T, Hindriks KV, Jonker CM, Yolum P (2013) Heuristic-based approaches for CP-nets in negotiation. In: Ito T, Zhang M, Robu V, Matsuo T (eds) Complex automated negotiations: Theories, models, and software competitions, Studies in computational intelligence, vol 435. Springer, Berlin, pp 113–123
36. Aydoğan R, Baarslag T, Hindriks KV, Jonker CM, Yolum P (2014) Heuristics for using CP-nets in utility-based negotiation without knowing utilities. Knowl Inf Syst 1–32
37. Boutilier C, Brafman RI, Domshlak C, Hoos HH, Poole D (2004) CP-nets: A tool for representing and reasoning with conditionalceteris paribus preference statements. J Artif Intell Res 21:135–191
38. Keeney RL, Raiffa H (1976) Decisions with multiple objectives. Cambridge University Press, Cambridge

39. Raiffa H, Richardson J, Metcalfe D (2003) Negotiation analysis: The science and art of collaborative decision making. Harvard University Press, Cambridge
40. Ito T, Zhang M, Robu V, Fatima S, Matsuo T (2011) New trends in agent-based complex automated negotiations, vol 383. Springer Science & Business Media, New York
41. Marsa-Maestre I, Lopez-Carmona MA, Ito T, Zhang M, Bai Q, Fujita K (2014) Novel insights in agent-based complex automated negotiation, vol 535. Springer, Japan
42. de Jonge D (2015) Negotiations over large agreement spaces. Ph.D. thesis, Universitat Autònoma de Barcelona
43. Ito T, Klein M, Hattori H (2008) A multi-issue negotiation protocol among agents with nonlinearutility functions. Multiagent Grid Syst 4(1):67–83
44. Klein M, Faratin P, Sayama H, Bar-Yam Y (2003) Negotiating complex contracts. Group Decis Negot 12:111–125
45. Lopez-Carmona MA, Marsa-Maestre I, Klein M, Ito T (2012) Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces. Auton Agents Multi-Agent Syst 24(3):485–535
46. Robu V, Somefun K, La Poutré JA (2005) Modeling complex multi-issue negotiations using utility graphs. Proceedings of the fourth international joint conference on autonomous agents and multiagent systems, AAMAS'05. ACM, New York, pp 280–287
47. Marsa-Maestre I, Lopez-Carmona MA, Velasco JR, Ito T, Klein V, Fujita K (2009) Balancing utility and deal probability for auction-based negotiations in highly nonlinear utility spaces. In: Proceedings of the 21st international joint conference on artifical intelligence, IJCAI'09. Morgan Kaufmann Publishers Inc., pp 214–219
48. Hadfi R, Ito T (2014) Addressing complexity in multi-issue negotiation via utility hypergraphs. In: Proceedings of the twenty-eighth AAAI conference on artificial intelligence
49. Chevaleyre Y, Endriss U, Estivie S, Maudet N (2004) Multiagent resource allocation with k-additive utility functions. In: Proceedings of the DIMACS-LAMSADE workshop on computerscience and decision theory, pp 83–100
50. Nisan N (2006) Bidding languages. Combinatorial auctions, Cambridge
51. Zeng D, Sycara KP (1997) Benefits of learning in negotiation. In: Proceedings of the fourteenth national conference on artificial intelligence and ninth conference on innovative applications of artificial intelligence, AAAI'97/IAAI'97. AAAI Press, pp 36–41
52. Zeng D, Sycara KP (1998) Bayesian learning in negotiation. Int J Hum Comput Stud 48(1):125–141
53. Carnevale PJD, Lawler EJ (1986) Time pressure and the development of integrative agreements in bilateral negotiations. J Conflict Resolut 30(4):636–659
54. Pruitt DG (1981) Negotiation behavior. Academic Press, New York
55. Kersten GE, Noronha SJ (1998) Rational agents, contract curves, and inefficient compromises. Trans Syst Man Cybern Part A 28(3):326–338
56. Hindriks KV, Tykhonov D (2010) Towards a quality assessment method for learning preference profiles in negotiation. In: Ketter W, La Poutré JA, Sadeh N, Shehory O, Walsh W (eds) Agent-Mediated electronic commerce and trading agent design and analysis. Lecture notes in business information processing, vol 44. Springer, Berlin, pp 46–59
57. Cheng C-B, Chan C-CH, Lin K-C (2006) Intelligent agents for e-marketplace: Negotiation with issue trade-offs by fuzzy inference systems. Decis Support Syst 42(2):626–638
58. Faratin P, Sierra C, Jennings NR (2002) Using similarity criteria to make issue trade-offs in automated negotiations. Artif Intell 142(2):205–237
59. Ito T, Hattori H, Klein M (2007) Multi-issue negotiation protocol for agents: exploring nonlinear utility spaces. Proceedings of the 20th international joint conference on artifical intelligence, IJCAI'07. Morgan Kaufmann Publishers Inc, San Francisco, pp 1347–1352
60. Jonker CM, Robu V, Treur J (2007) An agent architecture for multi-attribute negotiation using incomplete preference information. Auton Agents Multi-Agent Syst 15:221–252
61. Lin R, Oshrat Y, Kraus S (2009) Investigating the benefits of automated negotiations in enhancing people's negotiation skills. In: AAMAS'09: Proceedings of the 8th international conference on autonomous agents and multiagent systems. International foundation for autonomous agents and multiagent systems, Richland, pp 345–352

62. Karp AH, Wu R, Chen K-Y, Zhang A (2004) A game tree strategy for automated negotiation. In: Proceedings of the 5th ACM conference on electronic commerce, EC'04. ACM, New York, pp 228–229
63. Lin R, Kraus S, Wilkenfeld J, Barry J (2008) Negotiating with bounded rational agents in environments with incomplete information using an automated agent. Artif Intell 172(6–7):823–851
64. Kawaguchi S, Fujita K, Ito T (2012) Compromising strategy based on estimated maximum utility for automated negotiating agents. In: Ito T, Zhang M, Robu V, Fatima S, Matsuo T (eds) New trends in agent-based complex automated negotiations. Series of studies in computational intelligence, Springer, Berlin, pp 137–144
65. Kawaguchi S, Fujita K, Ito T (2013) AgentK2: Compromising strategy based on estimated maximum utility for automated negotiating agents. In: Ito T, Zhang M, Robu V, Matsuo T (eds) Complex automated negotiations: Theories, models, and software competitions. Studies in computational intelligence, vol 435. Springer, Berlin, pp 235–241
66. Williams CR (2012) Practical strategies for agent-based negotiation in complex environments. Ph.D. thesis, University of Southampton, 2012
67. Williams CR, Robu V, Gerding EH, Jennings NR (2012) Iamhaggler: A negotiation agent for complex environments. In: Ito T, Zhang M, Robu V, Fatima S, Matsuo T (eds) New trends in agent-based complex automated negotiations, Studies in computational intelligence. Springer, Berlin, pp 151–158
68. Williams CR, Robu V, Gerding EH, Jennings NR (2013) Iamhaggler 2011: A gaussian process regression based negotiation agent. In: Ito T, Zhang M, Robu V, Matsuo T (eds) Complex automated negotiations: theories, models, and software competitions, Studies in computational intelligence, vol 435. Springer, Berlin, pp 209–212
69. Hao J, Leung H-F (2012) ABiNeS: An adaptive bilateral negotiating strategy over multiple items. Proceedings of the The 2012 IEEE/WIC/ACM international joint conferences on web intelligence and intelligent agent technology, WI-IAT'12, vol 2. IEEE Computer Society, Washington, pp 95–102
70. Hao J, Leung H-F (2014) CUHK agent: An adaptive negotiation strategy for bilateral negotiations over multiple items. In: Marsa-Maestre I, Lopez-Carmona MA, Ito T, Zhang M, Bai Q, Fujita K (eds) Novel insights in agent-based complex automated negotiation, Studies in computational intelligence, vol 535. Springer, Japan, pp 171–179
71. Chen S, Ammar HB, Tuyls K, Weiss G (2013) Optimizing complex automated negotiation using sparse pseudo-input gaussian processes. In: Proceedings of the 2013 international conference on autonomous agents and multi-agent systems, AAMAS'13. International foundation for autonomous agents and multiagent systems, Richland, pp 704–714
72. Chen S, Weiss G (2012) An efficient and adaptive approach to negotiation in complex environments. In: De Raedt L, Bessiere C, Dubois D, Doherty P, Frasconi P, Heintz F, Lucas PJF (eds) ECAI, Frontiers in artificial intelligence and applications, vol 242. IOS Press, pp 228–233
73. Chen S, Weiss G (2012) A novel strategy for efficient negotiation in complex environments. In: Timm IJ, Guttmann C (eds) Multiagent system technologies, Lecture notes in computer science, vol 7598. Springer, Berlin, pp 68–82
74. Chen S, Weiss G (2013) An efficient automated negotiation strategy for complex environments. Eng Appl Artif Intell 26(10)
75. Ilany L, Gal Y(K) (2014) The simple-meta agent. In: Marsa-Maestre I, Lopez-Carmona MA, Ito T, Zhang M, Bai Q, Fujita K (eds) Novel insights in agent-based complex automated negotiation. Studies in computational intelligence, vol 535. Springer, Japan, pp 197–200
76. Ilany L, Gal Y (2015) Algorithm selection in bilateral negotiation. Auton Agents Multi-Agent Syst 1–27
77. Ilany L, Gal Y (2013) Algorithm selection in bilateral negotiation. In: Proceedings of the twenty-seventh AAAI conference on artificial intelligence (AAAI 2013)
78. Bartolini C, Preist C, Jennings NR (2002) A generic software framework for automated negotiation. In: First international conference on autonomous agent and multi-agent systems

79. Dumas M, Governatori G, Ter Hofstede AHM, Oaks P (2002) A formal approach to negotiating agents development. Electron Commer Res Appl 1(2):193–207
80. Ashri R, Rahwan I, Luck M (2003) Architectures for negotiating agents. In: Proceedings of the 3rd central and Eastern European conference on multi-agent systems. Springer, pp 136–146
81. Hindriks KV, Jonker CM, Tykhonov D (2008) Towards an open negotiation architecture for heterogeneous agents. In: Klusch M, Pechoucek M, Polleres A (eds) Cooperative information agents XII. Lecture notes in computer science, vol 5180. Springer, Berlin, pp 264–279
82. Matos N, Sierra C, Jennings NR (1998) Determining successful negotiation strategies: an evolutionary approach. In: Proceedings international conference on multi agent systems, pp 182–189
83. Eymann T (2001) Co-evolution of bargaining strategies in a decentralized multi-agent system. In: AAAI fall 2001 symposium on negotiation methods for autonomous cooperative systems, pp 126–134
84. Beam C, Segev A (1997) Automated negotiations: a survey of the state of the art. Wirtschaftsinformatik 39(3):263–268
85. Ros R, Sierra C (2006) A negotiation meta strategy combining trade-off and concession moves. Auton Agents Multi-Agent Syst 12:163–181
86. Lai H, Doong H-S, Kao C-C, Kersten GE (2006) Negotiators' communication, perception of their counterparts, and performance in dyadic e-negotiations. Group Decis Negot 15:429–447
87. Lomuscio AR, Wooldridge MJ, Jennings NR (2001) A classification scheme for negotiation in electronic commerce. In: Dignum F, Sierra C (eds) Agent mediated electronic commerce, Lecture notes in computer science, vol 1991. Springer, Berlin, pp 19–33
88. Thomas KW (1992) Conflict and conflict management: Reflections and update. J Organ Behav 13(3):265–274
89. Zachariassen F (2008) Negotiation strategies in supply chain management. Int J Phys Distrib Logistics Manage 38:764–781
90. An B, Sim KM, Tang LG, Miao CY, Shen ZQ, Cheng DJ (2008) Negotiation agents' decision making using markov chains. In: Ito T, Hattori H, Zhang M, Matsuo T (eds) Rational, robust, and secure negotiations in multi-agent systems, Studies in computational intelligence, vol 89. Springer, Berlin, pp 3–23
91. Lewicki RJ, Saunders DM, Barry B, Minton JW (2003) Essentials of negotiation. McGraw-Hill, Boston, MA
92. Gode DK, Sunder S (1993) Allocative efficiency in markets with zero intelligence (ZI) traders: Market as a partial substitute for individual rationality. J Polit Econ 101(1):119–137
93. Axelrod R (1984) The evolution of cooperation. Basic Books, New York
94. Axelrod R, Dion D (1998) The further evolution of cooperation. Science 242(4884):1385–1390
95. Baarslag T, Hindriks KV, Jonker CM (2013) A tit for tat negotiation strategy for real-time bilateral negotiations. In: Ito T, Zhang M, Robu V, Matsuo T (eds) Complex automated negotiations: theories, models, and software competitions, Studies in computational intelligence, vol 435. Springer, Berlin, Heidelberg, pp 229–233
96. Hindriks, Jonker CM, Tykhonov D (2009) The benefits of opponent models in negotiation. In: Proceedings of the 2009 IEEE/WIC/ACM international joint conference on web intelligence and intelligent agent technology, vol 2. IEEE Computer Society, pp 439–444
97. Rahwan I, Ramchurn S, Jennings NR, McBurney P, Parsons S, Sonenberg L (2003) Argumentation-based negotiation. Knowl Eng Rev 18(04):343–375
98. Sandholm T, Lesser VR (1996) Advantages of a leveled commitment contracting protocol. In: Clancey WJ, Weld DS (eds) Proceedings of the thirteenth national conference on artificial intelligence and eighth innovative applications of artificial intelligence conference, AAAI 96, IAAI 96, vol 1. AAAI Press/The MIT Press, Portland, Oregon, pp 126–133, 4–8 Aug 1996
99. Kolomvatsos K, Hadjieftymiades S (2014) On the use of particle swarm optimization and kernel density estimator in concurrent negotiations. Inf Sci 262:99–116
100. Williams CR, Robu V, Gerding EH, Jennings NR (2012) Towards a platform for concurrent negotiations in complex domain. In: Proceedings of the fifth international workshop on agent-based complex automated negotiations (ACAN 2012)

101. Williams CR, Robu V, Gerding EH, Jennings NR (2012) Negotiating concurrently with unknown opponents in complex, real-time domains. In: 20th European conference on artificial intelligence, vol 242, pp 834–839
102. Kolomvatsos K, Anagnostopoulos C, Hadjiefthymiades S (2013) Determining the optimal stopping time for automated negotiations. IEEE Trans Syst Man Cybern Syst 99:1–1
103. Leonardz B (1973) To stop or not to stop some elementary optimal stopping problems with economic interpretations. Almqvist & Wiksell, Stockholm
104. Mengxiao W, de Weerdt M, La Poutré JA (2013) Acceptance strategies for maximizing agent profits in online scheduling. In: David E, Robu V, Shehory O, Stein S, Symeonidis A (eds) Agent-Mediated electronic commerce. Designing trading strategies and mechanisms for electronic markets, Lecture notes in business information processing. vol 119. Springer, Berlin, Heidelberg, pp 115–128
105. van den Herik J, Donkers J, Spronck PHM (2005) Opponent modelling and commercial games. In: Kendall G, Lucas S (eds) Proceedings of the IEEE 2005 symposium on computational intelligence and games, pp 15–25
106. Baarslag T, Hendrikx MJC, Hindriks KV, Jonker CM (2015) Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. Auton Agents Multi-Agent Syst pp 1–50
107. Schadd FC, Bakkes S, Spronck PHM (2007) Opponent modeling in real-time strategy games. In: 8th International conference on intelligent games and simulation (GAME-ON 2007), pp 61–68
108. Hindriks KV, Jonker CM, Tykhonov D (2011) Let's dans! An analytic framework of negotiation dynamics and strategies. Web Intell Agent Syst 9(4):319–335
109. Rahman SA, Bahgat R, Farag GM (2011) Order statistics bayesian-mining agent modelling for automated negotiation. Informatica Int J Comput Inf 35(1):123–137
110. Jeonghwan Gwak and Kwang Mong Sim (2011) Bayesian learning based negotiation agents for supporting negotiation with incomplete information. Proceedings of the international multiconference of engineers and computer scientists 1:163–168
111. Haberland V, Miles S, Luck M (2012) Adaptive negotiation for resource intensive tasks in grids. In: STAIRS, pp 125–136
112. Jazayeriy H, Azmi-Murad M, Sulaiman N, Udizir N (2011) The learning of an opponent's approximate preferences in bilateral automated negotiation. J Theor Appl Electron Commer Res 6(3):65–84
113. Klos TB, Somefun K, La Poutré JA (2011) Automated interactive sales processes. IEEE Intell Syst 26(4):54–61
114. Oprea M (2002) An adaptive negotiation model for agent-based electronic commerce. Stud Inform Control 11(3):271–279
115. Oshrat Y, Lin R, Kraus S (2009) Facing the challenge of human-agent negotiations via effective general opponent modeling. Proceedings of the 8th international conference on autonomous agents and multiagent systems, AAMAS '09, vol 1. International foundation for autonomous agents and multiagent systems. Richland, SC, pp 377–384
116. Papaioannou IV, Roussaki IG, Anagnostou ME (2011) Multi-modal opponent behaviour prognosis in e-negotiations. In: Cabestany J, Rojas I, Joya G (eds) Advances in computational intelligence, vol 6691., Lecture notes in computer scienceSpringer, Berlin Heidelberg, pp 113–123
117. Papaioannou IV, Roussaki IG, Anagnostou ME (2008) Neural networks against genetic algorithms for negotiating agent behaviour prediction. Web Intell Agent Syst 6(2):217–233
118. Rau H, Tsai M-H, Chen C-W, Shiang W-J (2006) Learning-based automated negotiation between shipper and forwarder. Comput Ind Eng 51(3):464–481
119. Sim KM, Guo Y, Shi B (2009) BLGAN: Bayesian learning and genetic algorithm for supporting negotiation with incomplete information. IEEE Trans Syst Man Cybern Part B Cybern 39(1):198–211
120. Sim KM, Guo Y, Shi B (2007) Adaptive bargaining agents that negotiate optimally and rapidly. IEEE Congr Evol Comput IEEE, pp 1007–1014

121. Agrawal MK, Chari K (2009) Learning negotiation support systems in competitive negotiations: A study of negotiation behaviours and system impacts. Int J Intell Inf Technol 5(1):1–23
122. Brzostowski J, Kowalczyk R (2006) Predicting partner's behaviour in agent negotiation. Proceedings of the fifth international joint conference on autonomous agents and multiagent systems, AAMAS'06. ACM, New York, NY, USA, pp 355–361
123. Jakub Brzostowski J, Kowalczyk R (2006) Adaptive negotiation with on-line prediction of opponent behaviour in agent-based negotiations. Proceedings of the IEEE/WIC/ACM international conference on intelligent agent technology, IAT'06. IEEE Computer Society, Washington, DC, USA, pp 263–269
124. Fang F, Xin Y, Yun X, Haitao X (2008) An opponent's negotiation behavior model to facilitate buyer-seller negotiations in supply chain management. In: Electronic commerce security, international symposium, pp 582–587
125. Hou C (2004) Predicting agents tactics in automated negotiation. In: Proceedings of the IEEE/WIC/ACM international conference on intelligent agent technology. IEEE Computer Society, pp 127–133
126. Masvoula M, Halatsis C, Martakos D (2011) Predictive automated negotiators employing risk-seeking and risk-averse strategies. In: Iliadis L, Jayne C (eds) Engineering applications of neural network. IFIP Advances in information and communication technology, vol 363. Springer Boston, pp 325–334
127. Mudgal C, Vassileva J (2000) Bilateral negotiation with incomplete and uncertain information: A decision-theoretic approach using a model of the opponent. Proceedings of the 4th international workshop on cooperative information agents IV, the future of information agents in cyberspace, CIA'00. Springer, London, UK, pp 107–118
128. Narayanan V, Jennings NR (2006) Learning to negotiate optimally in non-stationary environments. In: Klusch M, Rovatsos M, Payne TR (eds) Cooperative information agents X, vol 4149., Lecture notes in computer scienceSpringer, Berlin, pp 288–300
129. Ozonat K, Singhal S (2010) Design of negotiation agents based on behavior models. In: Chen L, Triantafillou P, Suel T (eds) Web information systems engineering-WISE 2010, vol 6488., Lecture notes in computer scienceSpringer, Berlin, pp 308–321
130. Ren Z, Anumba CJ (2002) Learning in multi-agent systems: A case study of construction claims negotiation. Adv Eng Inf 16(4):265–275
131. Saha S, Biswas A, Sen S (2005) Modeling opponent decision in repeated one-shot negotiations. Proceedings of the fourth international joint conference on autonomous agents and multiagent systems, AAMAS'05. ACM, New York, NY, USA, pp 397–403
132. Williams CR, Robu V, Gerding EH, Jennings NR (2011) Using gaussian processes to optimise concession in complex negotiations against unknown opponents. In: Proceedings of the twenty-second international joint conference on artificial intelligence, IJCAI'11, vol 1. AAAI Press, pp 432–438
133. Chao Y, Ren F, Zhang M (2013) An adaptive bilateral negotiation model based on Bayesian learning. In: Ito T, Zhang M, Robu V, Matsuo T (eds) Complex automated negotiations: Theories, models, and software competitions, vol 435., Studies in computational intelligenceSpringer, Berlin Heidelberg, pp 75–93
134. Fatima S, Wooldridge MJ, Jennings NR (2004) Optimal negotiation of multiple issues in incomplete information settings. In: Proceedings of the third international joint conference on autonomous agents and multiagent systems, vol 3. IEEE Computer Society, pp 1080–1087
135. Baarslag T, Hendrikx MJC, Hindriks KV, Jonker CM (2012) Measuring the performance of online opponent models in automated bilateral negotiation. In: Thielscher M, Zhang D (eds) AI 2012: Advances in artificial intelligence, vol 7691., Lecture notes in computer scienceSpringer, Berlin, pp 1–14
136. Coehoorn RM, Jennings NR (2004) Learning an opponent's preferences to make effective multi-issue negotiation trade-offs. Proceedings of the 6th international conference on electronic commerce, ICEC'04. ACM, New York, NY, USA, pp 59–68
137. Lau RYK, Li Y, Song D, Kwok RC-W (2008) Knowledge discovery for adaptive negotiation agents in e-marketplaces. Decis Support Syst 45(2):310–323

138. Lin R, Kraus S, Wilkenfeld J, Barry J (2006) An automated agent for bilateral negotiation with bounded rational agents with incomplete information. Proceedings of the 2006 conference on ECAI 2006: 17th European conference on artificial intelligence. IOS Press, Amsterdam, The Netherlands, pp 270–274
139. Mok WWH, Sundarraj RP (2005) Learning algorithms for single-instance electronic negotiations using the time-dependent behavioral tactic. ACM Trans Internet Technol 5(1):195–230
140. Somefun K, La Poutré JA (2007) A fast method for learning non-linear preferences online using anonymous negotiation data. In: Fasli M, Shehory O (eds) Agent-Mediated electronic commerce. Automated negotiation and strategy design for electronic markets, Lecture notes in computer science, vol 4452. Springer, Berlin, pp 118–131
141. Zhang M, Tan Z, Zhao J, Li L (2008) A bayesian learning model in the agent-based bilateral negotiation between the coal producers and electric power generators. In: International symposium on intelligent information technology application workshops, IITAW '08, pp 859–862, Dec 2008
142. Baarslag T, Hendrikx MJC, Hindriks KV, Jonker CM (2013) Predicting the performance of opponent models in automated negotiation. In: International joint conferences on web intelligence (WI) and intelligent agent technologies (IAT), 2013 IEEE/WIC/ACM, vol 2, pp 59–66, Nov 2013
143. Bosse T, Jonker CM, van der Meij L, Robu V, Treur J (2005) A system for analysis of multi-issue negotiation. In: Unland R, Calisti M, Klusch M (eds) Software agent-based applications, platforms and development kits, whitestein series in software agent technologies. Birkhöuser Basel, pp 253–279
144. Buffett S, Comeau L, Spencer B, Fleming MW (2006) Detecting opponent concessions in multi-issue automated negotiation. Proceedings of the 8th international conference on electronic commerce: The new e-commerce: Innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet, ICEC'06. ACM, New York, NY, USA, pp 11–18
145. Frieder A, Miller G (2013) Value model agent: A novel preference profiler for negotiation with agents. In: Ito T, Zhang M, Robu V, Matsuo T (eds) Complex automated negotiations: Theories, models, and software competitions, Studies in computational intelligence, vol 435. Springer, Berlin, pp 199–203
146. Hindriks KV, Tykhonov D (2008) Opponent modelling in automated multi-issue negotiation using bayesian learning. Proceedings of the 7th international joint conference on autonomous agents and multiagent systems, AAMAS'08, vol 1. International foundation for autonomous agents and multiagent systems. Richland, SC, pp 331–338
147. Jonker CM, Robu V (2004) Automated multi-attribute negotiation with efficient use of incomplete preference information. Proceedings of the third international joint conference on autonomous agents and multiagent systems, AAMAS'04, IEEE Computer Society, vol 3. Washington, DC, USA, pp 1054–1061
148. Robu V, La Poutré JA (2006) Retrieving the structure of utility graphs used in multi-item negotiations through collaborative filtering of aggregate buyer preferences. In: Proceedings of the 2nd international workshop on rational, robust and secure negotiations in MAS. Springer
149. Saha S, Sen S (2005) A bayes net approach to argumentation based negotiation. In: Rahwan I, Moraïtis P, Reed C (eds) Argumentation in multi-agent systems, vol 3366., Lecture notes in computer scienceBerlin, Heidelberg, pp 208–222
150. Şerban LD, Silaghi GC, Litan CM (2012) AgentFSEGA-time constrained reasoning model for bilateral multi-issue negotiations. In: Ito T, Zhang M, Robu V, Fatima S, Matsuo T (eds) New trends in agent-based complex automated negotiations. Series of studies in computational intelligence, Springer, Berlin, Heidelberg, pp 159–165
151. van Galen Last N (2012) Agent Smith: Opponent model estimation in bilateral multi-issue negotiation. In: Ito T, Zhang M, Robu V, Fatima S, Matsuo T (eds) New trends in agent-based complex automated negotiations. Studies in computational intelligence, Springer, Berlin, Heidelberg, pp 167–174

152. van Krimpen T, Looije D, Hajizadeh S (2013) Hardheaded. In: Ito T, Zhang M, Robu V, Matsuo T (eds) Complex automated negotiations: Theories, models, and software competitions, Studies in computational intelligence, vol 435. Springer, Berlin, pp 223–227
153. Aydoğan R, Yolum P (2012) The effect of preference representation on learning preferences in negotiation. In: Ito T, Zhang M, Robu V, Fatima S, Matsuo T (eds) New trends in agent-based complex automated negotiations, Studies in computational intelligence, vol 383., Springer-Berlin, Heidelberg, pp 3–20
154. Aydoğan R, Yolum P (2006) Learning consumer preferences for content-oriented negotiation. In: AAMAS Workshop on business agents and the semantic web (BASeWEB). ACM, Press, pp 43–52, May 2006
155. Aydoğan R, Yolum P (2012) Learning opponent's preferences for effective negotiation: An approach based on concept learning. Auton Agents Multi-Agent Syst 24:104–140
156. Farag GM, AbdelRahman SES, Bahgat R, A-Moneim Atef M. Towards KDE mining approach for multi-agent negotiation. In: 2010 The 7th International conference on informatics and systems (INFOS). IEEE, pp 1–7, Mar 2010
157. Buffett S, Spencer B (2005) Learning opponents' preferences in multi-object automated negotiation. Proceedings of the 7th international conference on electronic commerce, ICEC '05. ACM, New York, NY, USA, pp 300–305
158. Buffett S, Spencer B (2007) A bayesian classifier for learning opponents' preferences in multi-object automated negotiation. Electron Commer Res Appl 6(3):274–284
159. Restificar A, Haddawy P (2004) Inferring implicit preferences from negotiation actions. In: International symposium on artificial intelligence and mathematics. Fort Lauderdale, Florida, USA, Jan 2004
160. Ren F, Zhang M (2007) Predicting partners' behaviors in negotiation by using regression analysis. In: Zhang Z, Siekmann J (eds) Knowledge science, engineering and management, vol 4798., Lecture notes in computer scienceSpringer, Berlin, Heidelberg, pp 165–176
161. Masvoula M (2013) Forecasting negotiation counterpart's offers: A focus on session-long learning agents. In: COGNITIVE 2013, the fifth international conference on advanced cognitive technologies and applications, pp 71–76
162. Carbonneau RA, Kersten GE, Vahidov RM (2008) Predicting opponent's moves in electronic negotiations using neural networks. Expert Syst Appl 34(2):1266–1273
163. Carbonneau RA, Kersten GE, Vahidov RM (2011) Pairwise issue modeling for negotiation counteroffer prediction using neural networks. Decis Support Syst 50(2):449–459
164. Lee CC, Ou-Yang C (2009) A neural networks approach for forecasting the supplier's bid prices in supplier selection negotiation process. Expert Syst Appl 36(2, Part 2):2961–2970
165. Lin R, Kraus S, Baarslag T, Tykhonov D, Hindriks KV, Jonker CM (2014) Genius: An integrated environment for supporting the design of generic automated negotiators. Comput Intell 30(1):48–70
166. Devaux L, Paraschiv C (2001) Bargaining on an internet agent-based market: Behavioral vs. optimizing agents. Electron Commer Res 1:371–401
167. Fatima SS, Wooldridge MJ, Jennings NR (2005) A comparative study of game theoretic and evolutionary models of bargaining for software agents. Artif Intell Rev 23(2):187–205
168. Henderson P, Crouch S, Walters RJ, Ni Q (2003) Comparison of some negotiation algorithms using a tournament-based approach. In: Carbonell JG, Siekmann J, Kowalczyk R, Müller JP, Tianfield H, Unland R (eds) Agent technologies, infrastructures, tools, and applications for E-Services, vol 2592. Lecture notes in computer science. springerBerlin, Heidelberg, pp 137–150
169. Manistersky E, Lin R, Kraus S (2008) Understanding how people design trading agents over time. In: Proceedings of AAMAS'08, pp. 1593–1596
170. Hindriks KV, Jonker CM, Kraus S, Lin R, Tykhonov D (2009) Genius: negotiation environment for heterogeneous agents. Proceedings of the 8th international conference on autonomous agents and multiagent systems, AAMAS'09, vol 2. International foundation for autonomous agents and multiagent systems. Richland, SC, pp 1397–1398

171. Lin R, Kraus S, Tykhonov D, Hindriks KV, Jonker CM (2011) Supporting the design of general automated negotiators. In: Proceedings of the second international workshop on agent-based complex automated negotiations (ACAN'09), vol 319. Springer, pp 69–87
172. Thiessen EM, Loucks DP, Stedinger JR (1998) Computer-assisted negotiations of water resources conflicts. GDN 7(2):109–129
173. Kersten GE, Noronha SJ (1999) WWW-based negotiation support: Design, implementation, and use. Decis Support Syst 25(2):135–154
174. Hales D (2002) Neg-o-net-a negotiation simulation test-bed. Technical Report CPM-02-109, CPM
175. Marsa-Maestre I, Klein M, de la Hoz E, Lopez-Carmona MA (2011) Negowiki: A set of community tools for the consistent comparison of negotiation approaches. In: Kinny D, Hsu JYJ, Governatori G, Ghose AK (eds) Agents in principle, agents in practice. Lecture notes in computer science, vol 7047. Springer, Berlin, Heidelberg, 424–435
176. Littman M, Zinkevich M (2006) The 2006 AAAI computer poker competition. ICGA J 29(3):166
177. Wellman MP, Wurman PR, O'Malley K, Bangera R, de Lin S, Reeves D, Walsh WE (2001) Designing the market game for a trading agent competition. IEEE Internet Comput 5(2):43–51
178. Greenwald A, Stone P (2001) Autonomous bidding agents in the trading agent competition. IEEE Internet Comput 5(2):52–60
179. Ketter W, Collins J, Reddy P, Flath C, de Weerdt M (2011) The power trading agent competition. ERIM report series reference No. ERS-2011-027-LIS
180. Niu J, Cai K, Parsons S, McBurney P, Gerding EH (2010) What the 2007 tac market design game tells us about effective auction mechanisms. Auton Agents Multi-Agent Syst 21:172–203
181. Stone P, Greenwald A (2005) The first international trading agent competition: autonomous bidding agents. Electron Commer Res 5(2):229–265
182. Da Costa AD, Lucena CJ, Da Silva VT, Azevedo SC, Soares FA (2008) Art competition: agent designs to handle negotiation challenges. Trust in agent societies: 11th international workshop, TRUST 2008, Estoril, Portugal, May 12–13, 2008. Revised selected and invited papers. Springer, Berlin, pp 244–272
183. Fullam KK, Klos TB, Muller G, Sabater J, Schlosser A, Barber KS, Rosenschein JS, Vercouter L, Voss M (2005) A specification of the agent reputation and trust (art) testbed: experimentation and competition for trust in agent societies. In: The 4th International joint conference on autonomous agents and multi-agent systems (AAMAS). ACM Press, pp 512–518
184. Baarslag T, Fujita K, Gerding EH, Hindriks KV, Ito T, Jennings NR, Jonker CM, Kraus S, Lin R, Robu V, Williams CR (2013) Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. Artif Intell 198:73–103
185. Baarslag T, Hindriks KV, Jonker CM, Kraus S, Lin R (2012) The first automated negotiating agents competition (ANAC 2010). In: Ito T, Zhang M, Robu V, Fatima S, Matsuo T (eds) New trends in agent-based complex automated negotiations. Studies in computational intelligence, vol 383. Springer, Berlin, Heidelberg, pp 113–135
186. Baarslag T, Hindriks KV, Hendrikx MJC, Dirkzwager ASY, Jonker CM (2014) Decoupling negotiating agents to explore the space of negotiation strategies. In: Marsa-Maestre I, Lopez-Carmona MA, Ito T, Zhang M, Bai Q, Fujita K (eds) Novel insights in agent-based complex automated negotiation. Studies in computational intelligence, vol 535. Springer, Japan, pp 61–83
187. Baarslag T, Hindriks KV, Jonker CM (2011) Towards a quantitative concession-based classification method of negotiation strategies. In: Kinny D, Yung-jen Hsu J, Governatori G, Ghose AK (eds) Agents in principle, agents in practice, lecture notes in computer science, vol 7047. Springer, Berlin, Heidelberg, pp 143–158
188. Baarslag T, Hindriks KV (2013) Accepting optimally in automated negotiation with incomplete information. In: Proceedings of the 2013 international conference on autonomous agents and multi-agent systems. AAMAS '13, Richland. International foundation for autonomous agents and multiagent systems, pp 715–722

http://www.springer.com/978-3-319-28242-8

Exploring the Strategy Space of Negotiating Agents
A Framework for Bidding, Learning and Accepting in Automated Negotiation
Baarslag, T.
2016, XXI, 276 p. 58 illus., 21 illus. in color., Hardcover
ISBN: 978-3-319-28242-8