Artif Intell Rev (2008) 29:1–44
DOI 10.1007/s10462-009-9107-8

Negotiation among autonomous computational agents: principles, analysis and challenges

Fernando Lopes · Michael Wooldridge · A. Q. Novais

Published online: 15 July 2009

Abstract Automated negotiation systems with software agents representing individuals or organizations and capable of reaching agreements through negotiation are becoming increasingly important and pervasive. Examples, to mention a few, include the industrial trend toward agent-based supply chain management, the business trend toward virtual enterprises, and the pivotal role that electronic commerce is increasingly assuming in many organizations. Artificial intelligence (AI) researchers have paid a great deal of attention to automated negotiation over the past decade and a number of prominent models have been proposed in the literature. These models exhibit fairly different features, make use of a diverse range of concepts, and show performance characteristics that vary significantly depending on the negotiation context. As a consequence, assessing and relating individual research contributions is a difficult task. Currently, there is a need to build a framework to define and characterize the essential features that are necessary to conduct automated negotiation and to compare the usage of key concepts in different publications. Furthermore, the development of such a framework can be an important step to identify the core elements of autonomous negotiating agents, to provide a coherent set of concepts related to automated negotiation, to assess progress in the field, and to highlight new research directions. Accordingly, this paper introduces a generic framework for automated negotiation. It describes, in detail, the components of the framework, assesses the sophistication of the majority of work in the AI literature on these components, and discusses a number of prominent models of negotiation. This paper also highlights some of the major challenges for future automated negotiation research.

F. Lopes (✉) · A. Q. Novais
Department of Modelling and Simulation, LNEG-National Research Institute,
Estrada do Paço do Lumiar 22, 1649-038 Lisbon, Portugal
e-mail: fernando.lopes@ineti.pt

A. Q. Novais
e-mail: augusto.novais@ineti.pt

M. Wooldridge
Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK
e-mail: M.J.Wooldridge@csc.liv.ac.uk

Keywords Automated negotiation · Negotiation framework · Pre-negotiation · Bargaining · Impasse · Renegotiation · Multi-agent systems · Autonomous agents · Negotiation systems

1 Introduction

Negotiation is an important form of social interaction—management and labor negotiate the terms of contracts, businesses negotiate to purchase raw materials and to sell products, diplomats of different nations negotiate peace accords, etc. Human negotiation is studied in the various branches of the social sciences, notably international relations, economics and social psychology (see e.g., Lewicki et al. 2003; Pruitt and Kim 2004; Thompson 2005). Automated negotiation is an active area of research in computer science in general and artificial intelligence in particular. The demands for systems composed of computational agents that are owned by different individuals or organizations and are capable of reaching agreements through negotiation are becoming increasingly important and pervasive. Examples, to mention a few, include:

1. the industrial trend toward agent-based supply chain management—agents representing business units or facilities negotiate the conditions for purchasing raw materials, decide and execute the scheduling, and negotiate the terms under which the products are delivered;
2. the business trend toward virtual enterprises—dynamic alliances of small, agile enterprises which together can take advantage of economies of scale when available;
3. the pivotal role that electronic commerce (e-commerce) is increasingly assuming in many organizations—e-commerce offers opportunities to significantly improve (make faster, cheaper, and more agile) the way that businesses interact with both customers and suppliers.

Artificial intelligence (AI) researchers have investigated the design of agents with negotiation competence from two main perspectives: a theoretical or formal mathematical perspective and a practical or system-building perspective. Researchers following the theoretical perspective have drawn heavily from game-theoretic and economic methods (see e.g., Rosenschein and Zlotkin 1994; Sandholm 1999; Kraus 2001; Fatima et al. 2005). Most theoretical models have some highly desirable properties, such as Pareto efficiency and the ability to guarantee convergence, but work with abstract problems under assumptions that often limit their applicability. On the other hand, researchers following the practical perspective have drawn heavily on social sciences techniques for understanding interaction and negotiation (see e.g., Müller 1996a; Jennings et al. 2001; Rahwan et al. 2004; Kersten and Lai 2007). Most computational models are being used successfully in a wide variety of real-world domains, but often without any rigorous theoretical underpinning (see next section for further discussion).

The various negotiation models that have been proposed in the literature exhibit fairly different features and make use of a diverse range of concepts, so assessing and relating individual research contributions is difficult. At present, there is a need to build a generic framework to define and characterize the essential features that are necessary to conduct automated negotiation and to compare the usage of key concepts in different publications. Furthermore, the development of such a framework can be an important step to identify the core elements of autonomous negotiating agents, to provide a coherent set of concepts related to automated negotiation, to understand the interrelationships of disparate research efforts, and to assess progress in the field. Also, such a framework can highlight new research

directions and raise new questions that exceed the capabilities of existing negotiation models, constituting an important tool for the development of more sophisticated negotiating agents. Arguably the lack of structure in the field leads to a sense of uncertainty about how to define the components or dimensions of such a framework.

Against this background, the purpose of this paper is threefold:

1. to examine the negotiation literature and to identify the most important components of a generic framework for automated negotiation;
2. to describe the various components of the framework and to establish a research agenda for automated negotiation;
3. to analyze and evaluate work in the AI literature according to each component of the framework and to highlight some of the major challenges for future research.

The paper is not meant as a survey of the field of automated negotiation. Rather, the description of the various components of the framework is generally undertaken with reference to work from artificial intelligence and various fields of the social sciences. For each component, an appraisal of the relative merits and drawbacks of the negotiation models that have been proposed in the literature is presented, together with a detailed description of a prominent model.

The remainder of this paper is structured as follows. Section 2 describes the main approaches followed by AI researchers for developing autonomous negotiating agents and summarizes their respective strengths and weaknesses. Additionally, Sect. 2 presents a multi-agent supply chain system to illustrate how negotiating agents operate in a real-world domain. Section 3 introduces a generic framework for automated negotiation. Sections 4–7 describe, in detail, the components of the framework, assess the sophistication of the majority of work in the AI literature on these components, and discuss a number of prominent models for negotiation. Section 8 presents concluding remarks and outlines some of the major challenges for future automated negotiation research. Finally, "Appendix" illustrates the applicability of the framework by classifying and comparing a representative sample of the most prominent models of negotiation that exist in the literature.

2 Autonomous agents and automated negotiation

2.1 Negotiation in artificial intelligence

AI researchers have paid a great deal of attention to automated negotiation over the last years. As noted earlier, some researchers have followed a theoretical or formal mathematical perspective and have mainly attempted to develop negotiation theories. To this end, they have drawn heavily on game-theoretic and economic methods. Most researchers have primarily focused on formal bargaining, auctions, market-oriented programming, contracting, and coalition formation. On the other hand, various researchers have followed a practical or system-building perspective and have mainly attempted to develop computer systems with negotiation competence. To this end, they have drawn heavily on social sciences techniques for understanding interaction and negotiation. Most researchers have primarily focused on the central process of moving toward agreement, notably the design and evaluation of negotiation protocols and negotiation strategies. Therefore, even though there exist several ways to classify existing models of negotiation, we adopt the following classification:

1. *theoretical models*: models for describing, specifying, and reasoning about the key features of negotiating agents;
2. *computational models*: models for specifying the key data structures of negotiating agents and the processes operating on these structures.

The theoretical models have some highly desirable properties such as maximizing social welfare, Pareto efficiency, stability, and the ability to guarantee convergence. However, most models are essentially static in the sense that they primarily focus on the outcome of negotiation rather than on the negotiation process itself—negotiation is mainly viewed as a process used to select a solution from a set of candidate solutions. Simply put, the negotiation process does not itself define the set of candidate solutions. Also, most models work with abstract problems and often fail to capture the richness of detail that would be necessary to successfully apply them in realistic domains. Furthermore, most models make the following restrictive assumptions: (1) the agents are rational, (2) the set of candidate solutions is fixed and known by all the agents, (3) each agent knows the other agents' potential payoffs for all candidate solutions, and (4) each agent knows the other agents' potential attitudes toward risk. These assumptions fail in most realistic environments due to the limited processing and communication capabilities of existing systems.

Most computational models are being used successfully in a wide variety of real-world domains. These models exhibit the following desirable features: (1) they focus on both the negotiation process and the negotiation outcome, (2) they are based on realistic assumptions, and (3) they make use of moderate computational resources to find acceptable solutions (according to the principles of bounded rationality, Simon 1981). However, most models have a number of limitations. Firstly, they often lack a rigorous theoretical underpinning—they are essentially ad hoc in nature. Secondly, they often lead to outcomes that are sub-optimal. Finally, there is often no precise understanding of how and why they behave the way they do. Consequently, they need extensive evaluation.

At this stage, we hasten to add four explanatory and cautionary notes. First, there are two fundamental approaches to the study of human negotiation: the game-theoretic approach (e.g., Osborne and Rubinstein 1990; Osborne 2004) and the behavioral approach (e.g., Pruitt and Carnevale 1993; Pruitt and Kim 2004). These two approaches have motivated the distinction between theoretical and computational models. Second, there are other ways to classify existing models of automated negotiation (see e.g., He et al. 2003). In particular, Jennings et al. (2001) present a slightly different classification that was adopted by various researchers (e.g., Rahwan et al. 2004). The authors discuss and analyze three classes of models:

1. *game-theoretic models*: provide clear analysis of specific negotiation situations and precise results concerning the optimal strategy every negotiator should choose, i.e., the strategy that maximizes negotiation outcome; negotiation proceeds by an iterative exchange of proposals and counterproposals;
2. *heuristic models*: provide general guidelines to assist negotiators and beneficial strategies for moving toward agreement, i.e., strategies that lead to good (rather than optimal) outcomes; negotiation also proceeds by an iterative exchange of proposals and counterproposals;
3. *argumentation-based models*: allow negotiators to argue about their mental attitudes during the negotiation process; negotiators can provide arguments to: (1) justify their negotiation stance, or (2) persuade other negotiators to change their negotiation stance; negotiation may proceed by an iterative exchange of proposals, counterproposals, threats, promises, etc.

The most obvious difference between our classification and Jennings et al.'s classification is that we consider only one class of models, referred to as computational models, for both heuristic and argumentation-based models. In fact, we acknowledge that negotiation proceeds through various phases, notably a beginning or initiation phase, a middle or problem-solving phase, and an ending or resolution phase (but see Sect. 3, below). In particular, the problem-solving phase seeks a solution for a dispute and is characterized by movement toward a mutually acceptable agreement —it often involves the exchange of offers and counter-offers, the communication of key pieces of information, and the submission of arguments. Accordingly, we do not consider distinct classes for models incorporating mechanisms based on either protocols of alternating offers (allowing agents to exchange offers and counter-offers) or more sophisticated protocols (allowing agents to provide feedback on the proposals they receive and/or arguments to support their negotiation stance). We believe our classification is rather natural and more suitable to the purpose of this paper. Third, the distinction between theoretical and computational models is not absolute and the reader may find some overlap between them. While some models make assumptions that limit their applicability, and thus, are clearly theoretical (e.g., Kraus et al. 1995), and other models make use of moderate computational resources to find acceptable solutions, and thus, are noticeably computational (e.g., Faratin et al. 1998), there is a degree of uncertainty in some cases (e.g., Sierra et al. 1998). This is caused by the fact that such models are based on realistic assumptions, but often make considerable demands on any implementation, mainly because they appeal to very rich representations of the agents and their environments. Finally, there is no universal technique or approach to automated negotiation that suits every problem domain. Nevertheless, the class of models referred to as computational models is gaining increasing popularity within the mainstream AI community and therefore will receive most of our attention in the remainder of this paper.

2.2 Agents for supply chain management

Multi-agent systems (MAS) are ideally suited to represent problems that have multiple problem solving entities and multiple problem solving methods (Jennings et al. 1998). The major motivations for the increasing interest in MAS research include the ability to solve problems in which data, expertise, or control is distributed, the ability to allow inter-operation of existing legacy systems, and the ability to enhance performance along the dimensions of computational efficiency, reliability, and robustness. Agent technology has been used to solve real-world problems in a range of industrial applications, including manufacturing, process control, telecommunications, air traffic control and transportation systems (see e.g., Weiss 1999; Wooldridge 2002; Bussmann et al. 2004; Fasli 2007).

Central to the design and effective operation of a multi-agent system are a core set of problems and research questions, notably (Sycara 1998):

1. *the design problem*: how to formulate, describe, decompose, and allocate different problems and synthesize results among a group of intelligent agents?
2. *the coordination problem*: how to ensure that agents act coherently, accommodating the local decisions or non-local effects and avoiding harmful interactions?

A supply chain is a network of facilities that performs the functions of procurement of raw materials from suppliers, transformation of these materials into intermediate goods and final products, and the delivery of these products to customers. A multi-agent supply chain system is a collection of autonomous computational agents, each responsible for one or more supply chain functions, and each interacting with other agents in the execution of their

responsibilities. Supply chain functions range from the ordering and receipt of raw materials to the distribution and delivery of final products, via the scheduling, production, and warehousing of intermediate goods and final products (Fox et al. 2000).

Now, the design problem consists mainly in distributing different supply chain functions across a number of autonomous agents. A typical distribution involves at least the following agents:

1. *sales agent*: responsible for acquiring orders from customers, negotiating with customers, and handling customer requests for modifying or canceling orders;
2. *logistics agent*: responsible for coordinating the plants and distribution centers of a manufacturing enterprise: it manages the movement of materials and products across the supply chain, from the suppliers of raw materials to the customers of finished goods;
3. *scheduling agent*: responsible for scheduling and rescheduling the activities of a manufacturing enterprise;
4. *resource management agent*: responsible for dynamically managing the availability of resources in order to execute the scheduled activities;
5. *supplier agents* and *customer agents*: the suppliers sell raw materials and the customers buy finished goods.

The coordination problem is focussed on ensuring that autonomous agents act in a tightly coordinated manner in order to effectively achieve their goals. This problem can be addressed, at least in part, by designing agents that are able to coordinate their activities through negotiation. Let us introduce a specific situation involving negotiation between the sales agent and the logistics agent.

> David, the director of Sales, has lined up two new orders for a total of 15,000 men's suits: one for 10,000 and the other for 5,000 men's suits. Martin, the director of Logistics, has already stated that it will take four months to make the suits. Together, they will gross over a million Euros, with a fine profit for the company. The problem is that Martin insists that the job will take four months and David's customer wants a 2-month turnaround. David and Martin are discussing and, so far, have accomplished little more than making each other angry. However, they can resolve their differences by negotiating a mutually beneficial agreement.

For illustrative purposes, we consider the negotiation process only from the viewpoint of David. There are four major issues of concern: quantity_1, date_1, quantity_2 and date_2. The first two issues are the most important to David due to the inherent customer demands—he wants fast action on the 10,000 suit order. Also, after a period of consultation with the customer, David concludes that he is firm about the 10,000 suit order (and is willing to wait up to five weeks for 10,000 suits, or ultimately 9,500 suits), but he is moderately soft about the 5,000 suit order (and is willing to wait up to six weeks for only 4,000 suits).

The opening stance and the pattern of concessions are two central elements of negotiation (Lewicki et al. 2003). Generally speaking, negotiators who demand too much will often fail to reach agreement and thereby do poorly. Those who demand too little will usually reach agreement but achieve low benefits (Pruitt and Carnevale 1993). For instance, if David is overly firm and keeps on demanding a two-month schedule, negotiation may break down. But if he is overly soft and concedes a lot, Martin will win.

Skilled negotiators are often moderately firm and hence between these two extremes. Furthermore, they often try to devise new alternatives by means of problem solving (Pruitt and Kim 2004). As noted, it is of higher priority for Sales to get fast action on the 10,000

suit order than the 5000 suit order. Suppose now that it is of higher priority for Logistics to avoid handling the 10,000 suit order. Also, Logistics prefers handling the minimal number of suits that Sales is willing to accept (4,000 suits in 6 weeks). These two departments have the makings of a logrolling deal—each party can yield on issues that are of low priority to itself and high priority to the other side. Accordingly, they can agree on the following superior solution: a 4-week schedule for 10,000 suits and a 6-week schedule for 4,000 suits.

3 A generic framework for automated negotiation

Negotiation is a discussion among conflicting parties with the aim of reaching agreement about a divergence of interest (Pruitt 1998). The list of situations that can be handled by negotiation is endless. Some situations are purely competitive, as when the parties have completely opposed interests. Other situations are purely cooperative, as when the parties have perfectly compatible interests. Most situations are mixed-motive, containing elements of both competitive and cooperative situations—the parties' interests are imperfectly correlated (Thompson 2005). There are, however, several characteristics common to most negotiation situations, including (Lewicki et al. 2003): (1) two or more parties, (2) a conflict among the parties—that is, what one party wants is not necessarily what the other parties want, and (3) an individual preference to search for agreement rather than to appeal to a higher authority, to permanently break off contact, or to fight openly. Hence conflict, while varying in importance and complexity, is present in most negotiation situations. Conflict also plays an important role in negotiation—it is the basis or the driving force of negotiation (Lewicki et al. 1999).

Negotiation, like other forms of social interaction, often proceeds through distinct phases or stages. A phase is a coherent period of interaction characterized by a dominant group of communicative acts that serves a set of related functions in the movement from initiation to a resolution of a dispute (Holmes 1992). Phase models provide a narrative explanation of the negotiation process, i.e., they identify sequences of events that constitute the story of negotiation (see e.g., Zartman and Berman 1982; Putnam et al. 1990; Greenhalgh 2001). Most models fit into a general structure of three phases: a beginning or initiation phase, a middle or problem-solving phase, and an ending or resolution phase. The initiation phase focuses on the preparation and planning for negotiation—it is marked by each party's efforts to emphasize points of difference and posture for positions. The problem-solving phase seeks a solution for a dispute—it is characterized by extensive interpersonal interaction, strategic maneuvers, and movement toward a mutually acceptable agreement. The resolution phase focuses on details and implementation of a final agreement. The parties often demand a gesture of good faith commitment to the agreement (close the deal) and determine who needs to do what once the documents are signed (implement the agreement). Several researchers point out, however, that a cautionary note is in order here—more research must be done about the nature of the phases themselves and about the nature of the processes that enable or drive the movement from phase to phase (Lewicki et al. 2003). Despite this, phase modelling offers much potential value in enhancing our understanding of negotiation.

The role of conflict as the driving force of negotiation and the three phases just outlined motivated the definition of the following groups of components or dimensions of a generic framework for automated negotiation:

Table 1 Components or dimensions of the negotiation framework

Group	Component or dimension
Preliminaries	1. Social conflict (detection and exploration)
	2. Negotiating parties (number of parties)
Pre-negotiation	3. Structuring of personal information (definition and execution of key pre-negotiation tasks)
	4. Analysis of the opponents (gathering and use of key information)
	5. Definition of the protocol and selection of the initial strategy
Actual negotiation	6. Exchange of offers and feedback information
	7. Argumentation (exchange of threats, promises, etc.)
	8. Learning (in negotiation)
	9. Dynamic strategic choice (selection of new strategies)
	10. Impasse resolution
Renegotiation	11. Analysis and improvement of the final agreement

1. preliminaries (the nature of negotiation);
2. pre-negotiation (preparing and planning for negotiation);
3. actual negotiation (moving toward agreement);
4. renegotiation (analyzing and improving the final agreement).

Further examination of the negotiation literature from the fields of social psychology (e.g., Pruitt 1981; Pruitt and Carnevale 1993; Van de Vliert 1997; Pruitt and Kim 2004), management science and economics (e.g., Fisher and Ury 1981; Raiffa 1982; Lewicki et al. 1999; Raiffa 2002; Lewicki et al. 2003; Thompson 2005), international relations (e.g., Snyder and Diesing 1977), and artificial intelligence (e.g., Müller 1996a; Sandholm 1999; Jennings et al. 2001; Lomuscio et al. 2003; Rahwan et al. 2004) motivated the definition of the dimensions in each group (see Table 1). For the sake of readability, a brief description of each dimension follows. A detailed description will be presented in Sects. 4–7, together with an assessment of the sophistication of the majority of work in the AI literature on each dimension. A number of prominent models of negotiation will also be analyzed.

The first two dimensions address the nature of negotiation. Specifically, the dimension "social conflict" accounts for the inevitability of conflict and the role conflict plays in driving negotiation—some models and systems may simplify or even ignore the inevitability of conflict, whereas others may recognize and explore conflict to drive negotiation. The dimension "negotiating parties" refers to the number of parties that can participate in negotiation—some models may deal with negotiation between two parties, whereas others may address the dynamics of multiparty negotiation.

The dimensions 3–5 address the operational and strategic process of preparing and planning for negotiation (referred to as pre-negotiation). In particular, the dimension "structuring personal information" refers to a number of key activities that every negotiator should attend to before starting to negotiate (e.g., definition and prioritization of the issues to negotiate)—some models may ignore such activities or simply address a few isolated activities,

whereas others may account for systematic preparation and planning for negotiation. The dimension "analysis of the opponent(s)" refers to the important step of gathering and using information about the opponent(s) to gain a clear sense of direction on how to proceed—some models may oversimplify or even ignore the importance of such information, whereas others may account for both a methodical acquisition and an efficient use of key pieces of information. The dimension "definition of the protocol and selection of the initial strategy" refers to two relevant pre-negotiation steps, namely joint agreement on a negotiation protocol and individual selection of a negotiation strategy—some models may simply consider pre-defined protocols and strategies, whereas others may explicitly address the definition of appropriate protocols and the selection of effective strategies for different negotiation situations.

The dimensions 6–10 address the central process of moving toward agreement (referred to as actual negotiation, or simply negotiation). Specifically, the dimension "exchange of offers and feedback information" refers to the communication of both offers and key pieces of information to permit a common definition of the negotiation situation to emerge—some models may focus on the communication of offers and oversimplify or even ignore the role of information in negotiation, whereas others may account for a richer form of interaction by integrating the communication of offers with the submission of feedback information. The dimension "argumentation" refers to the activities of structuring and presenting arguments to influence the attitude and behaviour of the negotiating parties—some models may ignore the role of arguments in negotiation, whereas others may support different types of arguments and account for their effective use throughout negotiation. The dimension "learning" refers to the acquisition of new knowledge and the ability to use that knowledge to improve negotiation performance—some models may not be able to learn, whereas others may incorporate learning mechanisms and exhibit richer (i.e., more effective) behaviours. The dimension "dynamic strategic choice" refers to the selection of different strategies as negotiation unfolds—some models may treat strategies as rigid or static elements that do not change during negotiation, whereas others may treat strategies as elements that do change throughout negotiation and incorporate effective approaches to dynamic strategic choice. The dimension "impasse resolution" accounts for the occurrence of impasses and their resolution—some models can ignore the occurrence of impasses, whereas others can manage "difficult" negotiations (that end in impasse) and incorporate effective approaches to impasse resolution.

The dimension 11 addresses the closing process of analyzing and improving the final agreement (referred to as renegotiation). Some models may oversimplify or even ignore the renegotiation process, whereas others may support the use of the final agreement as a plateau that can be revised and improved in a follow-up negotiation effort (i.e., may incorporate effective renegotiation approaches).

Now, there are other ways in which a negotiation framework can be viewed abstractly. Specifically, Wurman et al. (2001) focus on auctions and discuss the parameterization of the auction design space. The authors present an organizational framework that captures the essential features of most auction mechanisms and is organized along three axes, namely bidding rules (e.g., dominance rules, withdrawal and expiration rules, and activity rules), clearing policy (notably, a matching function), and information revelation policy (notably, a quote function). He et al. (2003) focus on agent-mediated electronic commerce, concentrating particularly on the business-to-consumer (B2C) and business-to-business (B2B) domains. The authors discuss and analyse the roles of autonomous agents in each domain, highlighting their use in a number of key activities, including product and merchant brokering, buyer coalition formation, partnership formation, and automated negotiation (notably,

auctions, bilateral negotiation, and contracting). Lomuscio et al. (2003) present a classification framework for automated negotiation in electronic commerce. The authors build on the work of Wurman et al. (2001) and define five main groups of parameters to classify negotiation mechanisms, namely cardinality of negotiation (e.g., number of parties and number of issues), characteristics of agents (e.g., role, rationality, social behaviour, and bidding strategy), characteristics of the environment (e.g., static or dynamic), event parameters (notably, parameters related to the specification of the negotiation protocol), information parameters (e.g., feedback information and arguments submitted by negotiators), and allocation parameters (e.g., parameters specific to many-to-one and many-to-many auctions).

Rahwan et al. (2004) make a distinction between autonomous agents with significant negotiation competence, referred to as classical negotiating agents (e.g., auction-based and bargaining agents), and argumentation-based negotiation (ABN) agents, i.e., agents capable of exchanging proposals and counterproposals and, in addition, able to generate and communicate different types of arguments. The authors present a conceptual framework defining the core elements and features of both the ABN agents (notably, argument and proposal generation, argument and proposal evaluation, and argument selection) and the interaction environment (namely, a common domain language, a communication language, a negotiation protocol, and information stores to keep externally accessible information). Bartolini et al. (2005) characterize the negotiation design space in terms of declarative rules. The authors present a conceptual framework built upon an abstract negotiation process (a generic process incorporating aspects of both human and automated negotiation), a taxonomy of negotiation rules (e.g., rules for admission, proposal validity, updating status, and termination), and an interaction protocol (a generic protocol specifying the flow of messages and involving three main phases, namely admission, proposal submission, and agreement formation). Additionally, they describe a prototype implementation of the framework.

Kersten and Lai (2007) focus on software systems to support and automate negotiation. The authors define four main types of software (namely, negotiation support systems, e-negotiation systems, e-negotiation tables, and negotiation software agents) and present three different system classifications (namely, classifications based on activeness, negotiation roles, and negotiation activities). Additionally, they discuss various system architectures (e.g., tightly coupled and loosely coupled architectures), several negotiation models (notably, models of the negotiation problem, the negotiator, and the negotiation process), and different software configurations (e.g., individual support and decentralized support configurations). Finally, they highlight the interdisciplinary nature of the studies conducted on negotiation software design, development, and deployment, and point out the need to build a research framework to study and compare different systems.

Overall, these and other existing frameworks primarily focus on the core elements of the agents engaged in negotiation and the essential features of the interaction environment that hosts these agents, placing emphasis on the middle or problem-solving phase of negotiation. Accordingly, we hasten to add the following observations. First, few frameworks explicitly acknowledge and characterize the role that conflict plays in driving agent interaction and negotiation. Second, few frameworks focus on the operational and strategic process of preparing and planning for negotiation, and thus, define and characterize the key activities that negotiators should attend to before actually starting to negotiate. Third, as noted, most frameworks focus on the central process of moving toward agreement. However, despite their power and elegance, important negotiation activities are still waiting to be discussed more thoroughly (e.g., understanding the nature of negotiations that are "difficult" to resolve

and can end in impasse, or analyzing effective approaches to impasse avoidance and/or resolution). Finally, few frameworks focus on the closing process of analyzing and improving the final agreement, and consequently, define and characterize the key activities needed to reopen a deal and try a post-settlement settlement.

This paper shifts the emphasis from the middle phase of negotiation to the role of conflict in negotiation and the three main phases of negotiation just outlined. It both extends and updates the previous material and also tries to present a more integrated and coherent view on the field of automated negotiation. Furthermore, it points to some research areas that seem quite underdeveloped and that offer exciting opportunities for future study (but see Sects. 4–7).

4 The nature of negotiation

4.1 Social conflict—driving force of negotiation

The term "conflict" originally meant a fight, battle, or struggle, that is, an overt confrontation between the parties. But its meaning has grown to include a sharp disagreement or opposition, as of interests or ideas. The term now has come to be so broadly applied that it is in danger of losing its status as a singular concept. A solution to this problem consists of adopting a restrictive meaning that builds on the source of conflict. Accordingly, conflict means a perceived divergence of interest—a perception that the parties' current aspirations are incompatible, i.e., if one party gets what it wants, the other (or others) will not be able to do so (Pruitt and Kim 2004).

Three levels of social conflict are commonly identified in the literature (see e.g., Lewicki et al. 2003):

1. interpersonal conflict (e.g., conflict between bosses and subordinates);
2. intra-group conflict (e.g., conflict among team and committee members);
3. inter-group conflict (e.g., conflict among nations).

Conflict at the interpersonal and group levels may exhibit fairly different characteristics. Nevertheless, conflict is present and constitutes the basis of most or all negotiation situations. It is not necessarily bad or good—it basically reveals a perceived divergence of interest.

Summary of the state of the art and exemplar negotiation model. Traditionally, AI researchers have viewed conflict as an "out-designed" concept (Klein 1991; Malsch and Weiss 2000) and have concentrated on the process of moving toward agreement. However, as multi-agent systems have become more complex and more demanding a new insight has gained ground in AI. Lander (1994) claimed that conflict is the focal point of interaction and must be explicitly addressed. Castelfranchi (2000) argued that conflict is important per se and must be accepted as normal social behavior. Wagner et al. (1999) pointed out that conflict in multi-agent systems is ubiquitous and, in some sense, all agent communication and interaction is motivated by the need to resolve conflict, by the need to deal with interdependence.

Various AI researchers have identified the source of conflict and explored the presence of conflict to drive negotiation. Von Martial (1992), for example, used negotiation to coordinate the plans of autonomous agents operating in cooperative environments. Positive plan relations (equality, favor, and subsumption) and negative plan relations (resource conflict and incompatibility) are detected using specialized algorithms. Plan coordination is done in a

centralized way by a coordinator agent using a negotiation protocol. Müller (1996b) developed a layered architecture for autonomous agents that incorporates methods for detecting and validating conflicts. Negotiation is triggered by real conflicts and involves mainly the selection of joint plans from a set of candidate plans and their subsequent transmission to the other parties (joint plans are plans that represent both actions of multiple agents and relationships among these actions). Lopes et al. (2002) presented a model for conflict detection and validation. The main components of the model are: (1) a library of conflict detection axioms, (2) a process for detecting potential conflicts, and (3) a process for validating potential conflicts. The authors pointed out that real conflicts raise negotiation problems and described how to define and generate a structure for these problems (but see Sect. 5.1).

Today there is a wide consensus that conflict is of particular importance to AI, but there is still a lack of theoretical understanding and important questions are still waiting to be addressed more thoroughly. We highlight the following:

1. How to develop effective methods for conflict detection?
2. How to acknowledge and explore the role of conflict as a driving force of negotiation?

Most AI researchers have paid little attention to these questions, notably to the role conflict pays in negotiation. Against this background, the remainder of this subsection discusses a piece of work that addresses some of the questions in a domain-independent way.

Lander (1994) and Lander and Lesser (1997) presented a multi-agent testbed framework for cooperative distributed search among heterogeneous agents (TEAM). The framework has three main components: (1) a common memory, (2) a manager (TMan), and (3) an agent set. The common memory is a standard partitioned blackboard memory. The manager has the primary functions of managing: (1) the agent set (e.g., addition and deletion of agents), (2) the common memory, and (3) the communication among agents and between TMan and agents (in the form of messages). The agent set comprises individual agents with specific capabilities allowing them to act as members of a team. The agents are benevolent (non-hostile) and are either globally cooperative (pursue a global goal) or locally cooperative (pursue local goals and interact only when doing so helps them to achieve their goals). They can have both representation heterogeneity (e.g., in the architecture) and knowledge heterogeneity (e.g., in the capabilities). The authors addressed distributed-search problems that have a natural sub-problem decomposition based on either knowledge or representation heterogeneity where each sub-problem is solved by an agent (examples include design and diagnosis). Distributed search involves four interwoven activities:

1. local search for a solution to some sub-problem;
2. integration of local sub-problem solutions into shared solutions;
3. information exchange to define and refine the shared search space of the agents;
4. assessment and reassessment of emerging solutions.

These activities are coordinated across a set of agents using distributed-search strategies. A distribute-search strategy specifies a global domain-independent and agent-independent algorithm for the application of a set of search operators. A distributed search operator implements a particular task at an agent. Negotiated search is the default strategy and exploits conflict to drive agent interaction and guide local search. Conflict resolution appears in different forms, notably extended-search (agents extend local search until founding a solution that does not conflict) and relaxation (agents expand local search by relaxing requirements of solutions). The authors claimed that their work acknowledges the inevitability of conflict among agents

and treats conflict as an integral part of problem solving (see also Sect. 6.1 and Table 2, in "Appendix", for a more in-depth discussion of this work).

4.2 Negotiating parties

Negotiation may involve two parties (bilateral negotiation) or more than two parties (multilateral negotiation) and one issue (single-issue negotiation) or many issues (multi-issue negotiation). Parties are the agents (or groups of agents with common interests) that participate in negotiation and issues are the resources to be allocated or the considerations to be resolved in negotiation (Thompson 2005). Multilateral negotiation may be classified according to the number of parties involved. One-to-many negotiation occurs when a single party negotiates with a number of other parties (e.g., auctions). For the purpose of analysis, this form of negotiation is often treated as a number of concurrent bilateral negotiations. Many-to-many negotiation occurs when multiple parties negotiate with many other parties (e.g., a task force deciding on a new product strategy for a manufacturing organization).

Multilateral negotiation is more complex, challenging, and difficult to manage than bilateral negotiation. There are five ways in which the complexity increases as three or more parties simultaneously engage in negotiation (Lewicki et al. 2003). First, there are simply more parties involved in the negotiation, which increases the number of different roles they may play and the demand for discussion time. Second, more parties bring more issues and positions to the table, and thus more perspectives must be presented and discussed. Third, negotiations become socially more complex—social norms often emerge that affect member participation, and there may be stronger pressures to conform and suppress disagreement. Fourth, negotiations become procedurally more complex, and the parties may have to negotiate a new process that allows them to coordinate their actions more effectively. Finally, negotiations become more strategically complex, because the parties must monitor the moves and actions of several other parties in determining what each will do next. Bazerman et al. (1988), Kramer (1991) and Thompson (2005) present useful analyses of the general characteristics and problems of multilateral negotiation.

Summary of the state of the art and exemplar negotiation model. The AI literature on negotiation is narrow in scope, although rich in detail. Most AI researchers have focused on bilateral negotiation—various negotiation models have been developed under the assumption that there are only two parties dealing directly with each other. There are also several pieces of work about one-to-many negotiation (e.g., Maes et al. 1999; Shintani et al. 2000; Zhang et al. 2005; Nguyen and Jennings 2005) and the remainder of this section describes a representative example. However, aside from the formal work on bargaining, auctions, market-oriented programming, contracting and coalition formation (see e.g., Sandholm 1999; Kraus 2001; Fatima et al. 2005), little work has addressed either negotiation involving more than two parties attempting to achieve a collective consensus or negotiation between teams or groups, i.e., one or more parties are comprised of a team of agents rather than a single agent. Yet multilateral negotiation is becoming increasingly important and pervasive—more than two parties, each with its own interests and positions, can work towards an effective agreement.

Li et al. (2004, 2006) addressed the problem of designing autonomous negotiating agents that take into consideration the effect of outside options (alternative candidates to negotiate an agreement) on negotiation behaviour. The authors considered three negotiation situations: (1) an agent wishing to purchase an item (buyer) engages in negotiation with other agent

capable of providing the desired item (seller), (2) a buyer engages in multiple concurrent bilateral negotiations with a set of available sellers (each bilateral negotiation is called a negotiation thread), and (3) a buyer engages in multiple concurrent bilateral negotiations with the available sellers and simultaneously takes into account uncertain sellers that may come dynamically in the future. They presented three negotiation models for the buyer agent:

1. a single-threaded model;
2. a synchronized multithreaded model;
3. a dynamic multithreaded model.

The single-threaded model does not consider outside options—a buyer and a seller negotiate the value of an item using a protocol of alternating offers (Osborne and Rubinstein 1990) and a set of time-dependent tactics (Faratin et al. 1998). The tactics define the pattern of concessions based on the reservation price, the negotiation deadline and, mainly, the amount of time that has elapsed since the beginning of negotiation. The synchronized multithreaded model builds on the single-threaded model and considers the presence of outside options—a buyer participates in multiple bilateral negotiations with different opponents. The negotiation behaviour of the buyer in a negotiation thread is influenced by the presence of the outside options that concurrently exist with that thread. Specifically, the buyer computes the reservation price in a thread from the value of the seller in that thread (the buyer's valuation of the item provided by the seller) and the reservation utility of the buyer in the thread (the expected utility of the buyer from all other threads). The utility that the buyer expects from a set of threads is estimated by heuristics based on either a reverse English auction or the outcome of a single-threaded negotiation with the opponent that gives more utility to the buyer. The dynamic multithreaded model expands the synchronized multithreaded model by considering future outside options (for an ongoing thread the outside options not only include the other simultaneous threads, but also the threads that may be launched in the future). The reservation utility of a thread is the expected utility of a multithreaded negotiation with stochastic thread number and uncertain item value. Table 2 summarizes the main features of this work (see "Appendix").

5 Pre-negotiation (preparing and planning for negotiation)

5.1 Structuring personal information

Successful human negotiators agree on one thing (Lewicki et al. 1999): the keys to success in negotiation are preparation and planning (usually referred to as pre-negotiation). Persuasive argumentation, skillful communication, clever maneuvering, and occasional histrionics may be important, but they cannot overcome the disadvantages created by a poor preparation and planning, nor can they help negotiators who have locked themselves into untenable positions before or during the early stages of negotiation. Although these procedures make the negotiation process interesting, the foundations for success are the preparation and planning that take place prior to actual negotiation.

Negotiators who carefully prepare and plan will make efforts to perform a number of activities, including (Raiffa 1982; Bazerman and Neale 1992; Lewicki et al. 2003):

1. defining the issues;
2. establishing the negotiating agenda;
3. prioritizing the issues;
4. defining the limits and targets.

First, every negotiator should define the issues to be discussed. In any negotiation, a detailed list of the issues at stake can be derived from a number of sources (e.g., an analysis of the conflict, or past experience in similar situations). Second, every negotiator should assemble all the issues that have been defined into a comprehensive list. The combination of lists from each side in the negotiation determines the negotiating agenda—the final set of issues to be deliberated. This task often involves interaction with the other parties—every negotiator discloses its list of issues in order to reach agreement about what will be discussed during actual negotiation. Third, every negotiator should prioritize the issues. Prioritization usually involves two steps: (1) determining which issues are most important and which are less important, and (2) determining whether the issues are connected or separate. Priorities can be set in a number of ways (e.g., to group the issues into categories of high, medium, and low importance, or to use standard techniques such as the nominal group technique). Priorities are often concealed, though in some situations every negotiator can disclose information about its priorities in order to find a mutually acceptable solution (but see Sects. 5.3 and 6.1).

The next step is to define two key points for each issue at stake in negotiation: (1) the limit (the point where every negotiator decides that it should stop the negotiation rather than to continue, because any settlement beyond this point is not minimally acceptable), and (2) the target point (the point where every negotiator realistically expects to achieve a settlement). The limit is also referred to as the resistance point, the reservation price or the walkaway point (Lewicki et al. 2003) and there are various ways to set it. A common method is to reason backward and thinking about what to do if a negotiation deal fell through (Thompson 2005). The target point is also referred to as the level of aspiration (Pruitt and Kim 2004) and there are also several ways to set it (e.g., to analyze both the current negotiation situation and the outcomes that other negotiators have achieved in similar situations). Negotiators are typically willing to disclose their target point, but work hard to conceal their limit.

Although it is important to distinguish among these activities, we hasten to add two cautionary notes. First, we assume that these activities can proceed linearly, in the order in which they were presented. Information often cannot be obtained and accumulated quite this simply and straightforwardly, however, and information about both the negotiation situation and the opposing negotiators discovered in performing some of the later activities may force to reconsider earlier activities. Second, the literature on negotiation discusses a number of different pre-negotiation templates, which tend to emphasize different activities in slightly different sequences (see e.g., Fisher and Ertel 1995; Greenhalgh 2001). Nevertheless, we have tried in this subsection to present the most important activities in the pre-negotiation process.

Summary of the state of the art and exemplar negotiation model. There has been little research on pre-negotiation activities. Most AI researchers have paid little attention to the preliminary activities of negotiation and, as stated, have focused on the process of moving toward agreement. Specifically, pre-negotiation research has been conducted in a slightly fragmented manner—rather than approaching the content of pre-negotiation plans as a coherent structure, researchers have focused on a few activities in isolation from other activities. Most researchers have developed models that address the formal representation of the issues to be deliberated, the formulation of their relative importance (the priorities and/or weights), and the definition of the limits. Zeng and Sycara (1998), for example, proposed a sequential decision-making model, called Bazaar, which is able to learn. The model formalizes the issues to negotiate, defines their reservation prices, and incorporates a negotiation

protocol, a set of negotiation strategies, and a learning mechanism. The negotiating agents keep histories of their interactions and update their beliefs, using Bayesian updating, after observing the environment and the behaviour of other agents (see also Sect. 6.3). Faratin et al. (1998), Faratin (2000) and Faratin et al. (2002) presented a model for bilateral service-oriented negotiation: an agent (the client) requires a service to be performed on its behalf by another agent (the server). The model incorporates a negotiation protocol, an issue manipulation protocol, three decision making mechanisms, and a meta-strategy mechanism. The negotiation protocol starts with a pre-negotiation phase aiming at formalizing the (pre-agreed) issues, their weights, and their reservation values (see also Sect. 6.4). Zhang et al. (2005) presented a model that handles multi-linked negotiation in multi-task, resource sharing environments (typically, one agent negotiates with multiple agents about different issues and the negotiation over one issue affects the negotiation over other issues). The model defines a multi-link negotiation problem, formalizes the features or attributes in negotiation, and incorporates a heuristic search algorithm to find a nearly optimal solution for that problem.

However, despite these and other relevant pieces of work, AI researchers have paid little attention to important pre-negotiation activities. In particular, little work has addressed the identification of the negotiation issues and the definition of the negotiation agenda (most researchers have considered pre-defined sets of issues and pre-agreed agendas). Also, there is little or no direct, verifiable evidence that systematic planning actually improves negotiation effectiveness. Against this background, the remainder of this subsection discusses a piece of work that addresses the majority of the pre-negotiation activities outlined above.

Lopes et al. (2002, 2004, 2005) addressed the problem of developing agents with competence for: (1) detecting and validating conflicts, and (2) resolving conflicts through negotiation. They presented a computationally tractable model that handles multi-issue negotiation. The main components of the model are:

1. a pre-negotiation model;
2. a negotiation protocol;
3. an individual model of the negotiation process;
4. a set of negotiation strategies and a set of negotiation tactics.

The pre-negotiation model formalizes the activities that each agent should attend to before starting to negotiate. These activities are the following: (1) definition of the negotiation problem, (2) identification and prioritization of the issues to negotiate, (3) formulation of the limits and aspirations, and (4) definition of the negotiation constraints. Specifically, the model defines the problem under negotiation, formalizes its structure (an And/Or tree), partially addresses the activity of identifying the negotiation issues (from the leaf nodes of the tree), and specifies the relative importance of the issues (their priorities and weights). Also, the model defines the limits, the targets, and the hard and soft constraints for the issues at stake. Hard constraints are non-relaxable constraints that specify threshold values for the issues. Soft constraints are relaxable constraints that specify minimum acceptable values for the issues. The second component of the negotiation model, i.e., the negotiation protocol formalizes the set of possible tasks that the agents can perform during the course of negotiation. The individual model of the negotiation process, the strategies, and the tactics formalize the tasks that each agent should perform in order to negotiate effectively. In particular, the strategies define the tactics to be used both at the beginning and during the course of negotiation. The tactics formalize the individual moves to be made at each point of negotiation. The main features of this work are summarized in Table 2 (see "Appendix").

5.2 Analyzing the opponent(s)

Effective planning requires information about the opposing negotiators and the ability to use that information to gain a clear sense of direction on how to proceed. The information which is most relevant depends on the opponent(s) and the negotiation situation. However, the following key pieces of data are often of great importance (Raiffa 1982; Roloff and Jordan 1992; Lewicki et al. 2003):

1. the intended limits and targets of the opponent(s);
2. the negotiating history of the opponent(s);
3. the intended strategies of the opponent(s).

Negotiators may speculate about the limits of the other parties. They can and often do think stereotypically, using their own initial interests as a guide, and assuming that others are like themselves and want similar things. In addition, negotiators may speculate about the targets of the other parties. However, they can also gather this data directly from them. To this end, they should exchange information about targets prior to actual negotiation and, if this is not possible, they should plan to collect that data during the opening stage of actual negotiation.

Negotiators should also gather information about the past behaviour of the other parties. How they have acted in the past is usually a good indicator of how they are likely to behave in the future. Therefore, a careful assessment of the negotiating history of the opponents can provide useful clues. To this end, negotiators can communicate to others who know the opposing negotiators, or to others who have been in their situation before. However, there is a potential danger in drawing conclusions from this information. Assuming that the opposing negotiators will act in the future as they have acted in the past is just an assumption—they can and do act differently in different circumstances at different times.

Finally, it would be most helpful to gain information about the intended strategies of the other parties. Although it is unlikely that they reveal any information about their strategies, negotiators can infer some information from whatever data they collect about them. Thus, interests and reputation may tell negotiators a great deal about what strategies the opposing negotiators intend to pursue.

In theory, it would be extremely useful to have as much information about the other parties as possible before opening the negotiation. Although it does take some time and effort to get some information, the results are usually more than worth the investment. In reality, however, it may not be possible to obtain any information through either direct communication or other research sources (e.g., asking others who know the opposing negotiators or have negotiated with them). If not, negotiators should attempt to see the negotiation from the perspective of the other parties and anticipate what they would want if they were negotiating from that point of view.

Summary of the state of the art and exemplar negotiation model. AI researchers have traditionally neglected the pre-negotiation step of gathering information about opposing negotiators, which we believe is unfortunate, because valuable information can often be obtained at little cost. Most researchers have expended little effort in approaching the content of pre-negotiation plans as a coherent structure and, as a result, they have not attempted to study both the acquisition of key pieces of information about the opponent(s) and the incorporation of such key pieces of information into these plans. Also, most researchers have not attempted to study the negotiation process from the perspective of the opponent(s). Some researchers have, however, developed models that consider (pre-collected and readily available) information about the opposing negotiators (typically encoded into probabilistic

distributions) and explore the use of that information to negotiate more effectively (see e.g., Zeng and Sycara 1998; Nguyen and Jennings 2005; Li et al. 2006). Despite the power and elegance of these pieces of work, the authors are aware of no work on explicitly modeling the pre-negotiation step of gathering information either directly or indirectly about the opponent(s). Therefore, the discussion of the aforementioned models is deferred to the next subsections.

5.3 Defining a protocol and selecting a strategy

Planning requires that negotiators agree on an appropriate protocol and be able to put together all the available information to select an effective strategy. The protocol specifies the rules that govern the interaction between the negotiating agents (Jennings et al. 2001). Specifically, the protocol defines the negotiation states (e.g., accepting proposals), the valid actions of the agents in particular states (e.g., which messages can be sent by whom, to whom, at what stage), and the events that cause negotiation states to change (e.g., proposal accepted). The strategy accounts for the individual decisions of each agent. While the protocol restricts the possible actions to perform, it often does not specify any particular action. Rather, it often marks branching points at which every agent has to make decisions according to its strategy.

The protocol can be more or less sophisticated depending on the type of information that can be exchanged between agents. Specifically, the protocol can be simple allowing agents to exchange only offers or proposals, i.e., solutions to the problem they face. Alternatively, the protocol can be richer, allowing agents to provide feedback on the proposals they receive. The feedback often takes the form of critiques, i.e., comments on which parts of proposals are acceptable or unacceptable. The protocol can also be very sophisticated, allowing agents to provide arguments to support their negotiation stance (e.g., to justify their stance or to persuade other agents to change their stance). These issues will be addressed in more detail in Sects. 6.1 and 6.2.

Negotiation strategies can reflect a variety of behaviours and lead to strikingly different outcomes. However, the following three fundamental groups of strategies are commonly used by human negotiators (Van de Vliert 1997; Pruitt and Kim 2004):

1. *contending* (also called *competing* or *dominating*): negotiators who employ strategies in this group maintain their aspirations and try to persuade or force the opposing negotiators to yield;
2. *concession making* (also called *yielding*, *accommodating* or *obliging*): negotiators reduce their aspirations to partially or totally accommodate the opposing negotiators; they work toward *compromise agreements*; such agreements are achieved when the parties concede to a middle ground on some obvious dimension or dimensions that link their initial proposals (Pruitt 1981);
3. *problem solving* (also called *collaborating* or *integrating*): negotiators maintain their aspirations and try to find ways of reconciling them with the aspirations of the opposing negotiators; they work toward *integrative agreements*, i.e., agreements that provide higher joint benefit than compromise agreements (Carnevale and Pruitt 1992).

Although it is important to distinguish among these three groups of strategies, we hasten to add several cautionary notes. First, we consider that the strategies in these groups are appropriate both in bilateral and multilateral negotiation. Most theory and research in this area focuses on negotiation between two parties, though some researchers have investigated the behaviour and effectiveness of various strategies in multiparty contexts (see e.g., Brett 1991; Weingart et al. 1993; Thompson 2005). Second, most strategies are implemented through

a wide variety of tactics. The line between strategies and tactics often seems indistinct, but one major difference is that of scope. Tactics are short-term moves designed to enact broad (or high-level) strategies (Lewicki et al. 2003). Third, most negotiation situations call forth a combination, and often a sequence, of strategies from different groups. Rarely is a strategy of one group used to the exclusion of the strategies of the other groups. Therefore, we address in this section the selection of an initial strategy (Sect. 6.4 discusses the change of strategy as negotiation unfolds). Finally, concession making strategies are essentially unilateral strategies. By contrast, contending and problem solving strategies work as intended only when they are accompanied by effective social influence. Hence, the strategies in these two groups (and the tactics that enact them) will receive the bulk of our attention in this subsection.

Contentious or competitive behaviour aims at imposing preferred solutions on the opposing negotiators. This behaviour can take several forms, including (Pruitt 1981; Pruitt and Carnevale 1993):

1. appearing firm;
2. reducing the resistance of the others to concession making;
3. imposing time pressure.

Appearing firm can be imposed by making high initial demands, conceding slowly, and making statements or moves to underscore the determination to hold firm at a particular offer (positional commitments). Efforts to reduce the resistance of others to concession making include persuasive arguments, threats, and promises. Time pressure can be imposed by announcing a deadline, threatening to turn to another partner, imposing continuing costs on the others, and dragging out the negotiation. Research indicates that contentious behavior, if successful, enhances one's own outcomes while diminishing those of the other party (Pruitt 1998). However, there are three pitfalls associated with this behavior. One is that it tends to crowd out problem solving, thereby reducing the likelihood of finding integrative agreements (Weingart et al. 1990). The second is that, when limits are high, it often leads to failure to reach agreement (Zubek et al. 1992). The third is that it is often imitated by the other party for defensive reasons—positional commitments encourage counter-commitments, threats elicit counter-threats and, in general, arguments produce counter-arguments. Such imitation may lead to a conflict spiral that produces serious escalation (Pruitt and Kim 2004).

Problem solving or collaborative behaviour aims at finding agreements that appeal to all sides (integrative agreements). The host of existing problem solving strategies includes (Pruitt 1981; Fisher and Ury 1981):

1. *expanding the "pie"*: negotiators increase the available resources in a way that all sides can achieve their goals;
2. *nonspecific compensation*: one party achieves its goals and pays off the opponent(s) for accommodating its interests;
3. *logrolling*: two parties agree to trade off among the issues under consideration so that each party concedes on the issues that are of low priority to itself and high priority to the other party;
4. *cost cutting*: one party achieves its goals and the costs (of successful goal achievement) imposed on the opponent(s) are reduced or eliminated;
5. *bridging*: negotiators search for novel alternatives that satisfy the interests (goals, values, needs) underlying their overt positions.

These strategies are implemented by different sets of tactics and require progressively more information about the other parties (but see Sect. 6.1). Integrative agreements are often advantageous to the parties, both individually and collectively, because they are mutually

rewarding, they are more likely to be complied with, they can allow the parties to avoid potential stalemates, and they can foster harmonious relations between the parties (Pruitt and Carnevale 1993). The first and third advantages have been supported by research (Pruitt 1981; Zubek et al. 1992).

The selection of an appropriate initial strategy is a critical step in preparing for negotiation. Pruitt and Rubin (1986) and Pruitt and Carnevale (1993) proposed the dual concern model for strategic choice in bilateral negotiation. The model asserts that negotiation strategies result from the conjunction of self-concern (concern about own outcomes) and other-concern (concern about the other party's outcomes). It makes the following predictions: (1) high self-concern coupled with low other-concern encourages contending, (2) high other-concern and low self-concern encourages concession making, (3) high self-concern and high other-concern encourages problem solving, and (4) low self-concern and low other-concern encourages inaction. The dual concern model has received strong support from various studies (De Dreu et al. 2000). Savage et al. (1989) proposed a similar model of strategic choice, which states that negotiation strategies result from the interplay of concern about own outcomes and concern about the relationship with the other party. The model makes similar predictions about the choice of negotiation strategies. Lewicki et al. (2003) discussed the factors that affect how negotiation strategies are designed and presented a general model of strategic choice. They suggested that negotiation strategies result from the interplay of several factors, mainly targets, trust and openness, context, and outcomes. Pruitt and Kim (2004) discussed the dual concern model and various other approaches to strategic choice, notably the perceived feasibility approach, which traces strategic choice to the perceived likelihood of success and the cost or risk of enacting the various strategies. A detailed discussion of this approach is beyond the scope of this paper.

Summary of the state of the art and exemplar negotiation model. Traditionally, AI researchers have concentrated on the formalization of negotiation protocols and negotiation strategies. Most researchers have developed models that include specific protocols (notably, the alternating offers protocol) and libraries of negotiation strategies (notably, concession and problem solving strategies). They have investigated the behaviour of these strategies to determine the most effective in various negotiation situations. Sathi and Fox (1989), for example, investigated constraint-directed negotiation in the domain of resource allocation in an engineering organization. They described and evaluated three constraint satisfaction and relaxation operators that were directly inspired by three strategies developed in the realm of human relations, namely logrolling, bridging and unlinking. Faratin et al. (1998), Faratin (2000) and Faratin et al. (2002) presented a model for bilateral service-oriented negotiation that defines a range of strategies (functions that map a matrix of real numbers ranging from zero to one into another similar matrix) and three groups of concession tactics: time dependent (functions of time), resource dependent (functions of limited resources), and behavior dependent or imitative (functions of the opponent's behavior). The model also defines a linear algorithm for making issue trade-offs (concessions on some issues and higher demands on other issues). The authors performed two sets of experiments to empirically evaluate these components of the model in various types of negotiation situations (see also Sect. 6.4). Lopes et al. (2002, 2004, 2005) presented a negotiation model that formalizes various problem solving and concession strategies commonly used by human negotiators. The authors described two sets of experiments conducted to: (1) investigate the behavior of the strategies and their associated tactics, and (2) evaluate the effect of the strategies and tactics both on the process and the outcome of negotiation. The experimental results confirmed a number of well-documented conclusions about human negotiation (this work was discussed in-depth in Sect. 5.1).

However, despite these and other relevant efforts to define negotiation strategies and to evaluate their effectiveness, AI researchers have paid little attention to a number of issues related to strategic choice. We highlight the following:

1. Which are the significant characteristics of both the negotiating parties and the negotiation situations that influence the selection of negotiation strategies?
2. How to define a taxonomy of such characteristics that would guide a system developer (or the agents on their behalf) in selecting effective negotiation strategies? Likewise, how to define negotiation strategies that include specifications of the characteristics for which they are more suitable?

There have been few attempts to develop models that incorporate effective approaches to strategic choice. The remainder of this section presents a piece of work that accounts for a simple, albeit effective, form of strategy selection.

Nguyen and Jennings (2003, 2004, 2005) presented a model that handles one-to-many negotiation in service-oriented contexts. The authors considered a particular negotiation situation where an agent wishing to purchase a service (buyer) engages in multiple concurrent bilateral negotiations with a set of agents capable of providing that service (sellers). The main components of the buyer are:

1. a coordinator;
2. a number of negotiation threads (one per seller);
3. a commitment manager.

The coordinator decides the negotiation strategy for each thread using the following key pieces of data:

1. a probability distribution over the types of the sellers;
2. a percentage of success matrix;
3. a pay off matrix.

The sellers can be of two types: conceder (willing to concede in a search for deals) or non-conceder (willing to adopt a tough stance). The percentage of success matrix measures the chance of having an agreement as the outcome of negotiation when the buyer applies a particular strategy to negotiate with a specific type of seller. The pay off matrix measures the average utility value of the agreement reached in similar situations. For each thread, the coordinator calculates the expected utility of applying different strategies to negotiate with a particular seller and selects the strategy that maximizes the expected utility. The following groups of strategies (and tactics) are available to the coordinator (Faratin et al. 1998): conceder, linear and tough. After finishing with the first seller (a randomly picked agent), the coordinator uses a Bayesian mechanism to update the probability distribution of the agent types and continues on with the second and subsequent sellers. The coordinator is also responsible for dynamically changing the strategies and for coordinating all the negotiation threads. Each thread interacts with a particular seller according to the alternating offers protocol (Osborne and Rubinstein 1990) and is mainly responsible for making offers and for deciding whether or not to accept an offer. The various threads mutually influence one another, i.e., the progress and agreement in one thread is used to influence the behaviour of the other ongoing threads. Specifically, each thread inherits various global parameters of the buyer agent, including the reservation value and the negotiation deadline. When a thread reaches a deal with a particular seller, the coordinator notifies all the other threads of the new reservation value (the utility of the deal) and may change the strategy for some of them. The third and last component of the buyer is the commitment manager, which is responsible for handling any issue that is related to commitment and decommitment (see also Sect. 7.1 and Table 2, in "Appendix").

6 Actual negotiation (moving toward agreement)

6.1 Offers and feedback information

The heart of negotiation is the exchange of offers and counter-offers (Raiffa 1982; Tutzauer 1992). In "good faith" negotiation, offers are made, and are either accepted or returned with counter-offers. There is an unstated assumption that the parties will show their commitment to the process of finding a solution by making concessions, and not simply by rejecting the offers of the others out of hand. To do so is often seen as "bad faith" bargaining. Failure to make concessions also conveys an image of hardline bargaining and unwillingness to negotiate. Hopefully, through the exchange of offers and counter-offers a point is reached on which the parties will agree.

Negotiation involves, however, more than a series of proposed settlements—a great deal of information is often conveyed permitting a "common definition" of the situation to emerge (Lewicki et al. 2003). The exchange of information is closely related to two dilemmas that all negotiators face (identified by Kelley 1966). The first dilemma, the dilemma of honesty, concerns how much of the truth to tell to the other parties. On the one hand, if negotiators tell everything about the current situation, then the other parties may take advantage of it. On the other hand, if negotiators do not tell anything about the current situation, then negotiation may evolve to a stalemate. The second dilemma, the dilemma of trust, concerns how much to believe of what the other parties tell to negotiators. To what extent negotiators should trust the other parties depend on a number of factors, including how they have negotiated in the past, their actual negotiation reputation, and the present circumstances.

Now, the decision about what information to communicate depends to some extent to the attitude of the negotiators. Competitive negotiators often seek their own advantage, mainly through gathering information about the other parties, concealing information about them, attempting to mislead, and using manipulative actions. In contrast, collaborative negotiators often create the conditions for effective information exchange, seek insight into the goals and interests of the other parties, and frequently use the five problem solving strategies outlined in the previous section. These strategies require progressively more information about the opponent(s) (Pruitt and Kim 2004). Specifically, the information requirements for expanding the "pie" are the demands of the opponent(s). Information about both the realms of value to the opponent(s) and how badly they are hurt by making concessions are useful to devise solutions by nonspecific compensation. Information about the priorities among the issues under discussion is helpful for developing solutions by logrolling. Cost cutting requires a deeper kind of information—it involves knowing something about the interests underlying the positions of the opponent(s). Bridging requires information about both the nature of the interests of all the parties and the priorities among the various interests. Ample research evidence indicates that effective information exchange promotes the development of good integrative solutions (see e.g., Thompson 1991; Olekalns et al. 1996; Butler 1999; Thompson 2005).

Summary of the state of the art and exemplar negotiation model. Traditionally, AI researchers have concentrated on the offer/counter-offer process—they have developed models that support mainly the generation and communication of offers and counter-offers. However, a number of researchers have also recognized and explored the importance of information in negotiation. Sycara (1987), for example, developed the Persuader system for resolving goal conflicts in the domain of labor relations. The system can perform three main tasks: (1) generation of proposals, (2) modification of rejected proposals based on feedback from

a dissenting party, and (3) persuasive argumentation. The modification of rejected proposals involves asking a rejecting party for feedback concerning the objectionable issues, the importance of these issues, and the reason for rejection, and is done through a process called re-analysis (see also Sect. 6.5). Garrido and Sycara (1996) presented a multi-agent meeting scheduling system where the agents can exchange proposals (values for date, start-time, and duration of meetings) and local information (preference values of time intervals). They found that revealing meeting preferences results in superior meeting schedule performance (in terms of efficiency and meeting quality). Lopes et al. (2002, 2004, 2005) presented a model that supports the communication of simple statements about relevant aspects of the negotiation process (e.g., the priorities of the issues under consideration). The authors used that information to formalize various problem solving strategies, notably logrolling. As stated earlier, they performed several experiments to empirically evaluate the effectiveness of these and other strategies in different negotiation situations (recall the exemplar model of Sect. 5.1).

These and other pieces of work may be viewed as attempts to integrate the communication of offers with the communication of feedback information. This information often takes the form of comments on which parts of the offers the agents like or dislike (e.g., explicit constraints on the issues). From such information, the agents are in a position to generate proposals that are more likely to be acceptable by their opponents. This provides a richer description of the negotiation process, one in which the offers both influence and are influenced by the complete communication situation. However, little attention has been given to a number of important issues, notably:

1. How to define and communicate key pieces of information that promote the development of integrative solutions (e.g., the nature and priorities of the agents' interests)? How to use that information to define important problem solving strategies (e.g., cost cutting and bridging)?
2. Under what conditions does the communication of feedback information effectively improve negotiation performance? How to evaluate the effect of communicating key pieces of information both on the convergence of the negotiation process and on the outcome of negotiation?

A cautionary and explanatory note is in order here. We make a distinction between feedback information and arguments. Although it is conceptually useful to distinguish between these two research topics, we acknowledge the existence of some overlap between them. A number of researchers go further and consider feedback information a particular type of argument. Jung et al. (2001), for example, presented a model for collaborative agents that attempt to solve a distributed constraint satisfaction problem (DCSP). The authors extended a standard DCSP algorithm (Yokoo and Hirayama 1998) by allowing agents to exchange both the values of the local variables and the local constraints (or arguments, as they called them). Parsons and Jennings (1996), Parsons et al. (1998) and Parsons and McBurney (2003) considered that all communication is argument-based, i.e., they only permit agents to make assertions that are backed by arguments (see Sect. 6.2 for a more in-depth discussion). Nevertheless, we have tried in this subsection to discuss some important pieces of work that support both the exchange of offers and the exchange of feedback information permitting a common definition of the situation to emerge (the remainder of the subsection presents a detailed description of a representative example). Section 6.2 discusses the role of arguments in negotiation placing emphasis on threats, promises, and persuasive arguments.

Lander (1994) and Lander and Lesser (1997) presented a multi-agent testbed framework that supports cooperative distributed search among reusable heterogeneous agents (recall the exemplar model of Sect. 4.1). Distributed search involves four interwoven activities: (1) local

search for a solution to some sub-problem, (2) integration of local sub-problem solutions into shared solutions, (3) information exchange to define and refine the shared search space of the agents, and (4) assessment and reassessment of emerging solutions. The authors described a distributed search strategy, called negotiated search, which incorporates some basic types of conflict management (notably, extended-search and requirement-relaxation). Negotiated search is basically a multipath incremental-extension algorithm, meaning that:

1. each solution is initiated by a single agent during some system processing cycle and then evaluated and extended by other agents during future processing cycles;
2. multiple solutions are constructed simultaneously.

Specifically, search is initiated by a problem specification that defines the form of a solution. This specification is further placed in the common memory and some agents use constraining information from the specification and their local solution requirements to propose one or more partial solutions. These solutions are then evaluated and extended by other agents. When a particular solution cannot be extended by some agent due to a conflict with existing solution attributes, there are two possible outcomes: (1) if the conflict is caused by violation of some hard (non-relaxable) constraint the solution path is pruned, or (2) if the conflict is caused by the violation of some soft (relaxable) constraint, the solution is saved and viewed as a potential compromise. In the first case, no more work will be done on the solution, and, to the extent that the violated constraint can be communicated to other agents, future counter-proposals will not violate the same constraint. In the second case, the violated constraint may eventually be relaxed and, if that happens, the potential compromise will become a viable solution again. In either of the two cases, conflict is used as the trigger for the communication of feedback information (i.e., local constraining information). The authors performed a set of experiments to empirically demonstrate the effectiveness of sharing potentially useful information among agents during distributed search. They revealed that information sharing can positively affect both solution quality (in terms of the monetary cost) and runtime (in terms of the elapsed real time from the invocation to the termination of the system). The main features of this work are summarized in Table 2, in "Appendix".

6.2 Persuasive arguments, threats and promises

Successful argumentation is an important component of effective negotiation. While negotiators may spend a great deal of time in assembling information to support their viewpoints, structuring and presenting effective arguments and counter-arguments are also powerful aspects of negotiation. Some researchers go further and consider that the core of negotiation is reciprocal offer and counter offer, argument and counter-argument in an attempt to agree upon outcomes mutually perceived as beneficial. Other researchers even consider that most of what happens in negotiation is the assertion of arguments by one side, and the response with counter-arguments by the other parties (Keough 1992).

Persuasive arguments are often used by a negotiator, the persuader, as a means to convince another negotiator, the target (or persuadee), to accept a particular proposal. The nature and types of persuasive arguments can vary enormously. However, the following types are commonly though to have persuasive force in real-life negotiations (Sycara 1987; Kraus et al. 1998):

1. *appeal to precedents as counterexamples*: to convey to the target a contradiction between what it demands and past actions;
2. *appeal to prevailing practice*: to convey to the target that a proposal will further its goals since it has furthered others' goals in the past;

3. *appeal to self-interest*: to convince the target that accepting a proposal will enable achievement of a high-importance goal.

These three argument types are not meant to constitute a typology of persuasive arguments. Indeed, it has been pointed out that it is not possible to present such an authoritative classification, since arguments must be interpreted and are effective within a particular context and domain (Toulmin et al. 1979). The effectiveness of these and other types of persuasive arguments has been a topic of extensive research by social sciences researchers. A detailed review of the literature is beyond the scope of this paper (see however,Chaiken et al. 2000).

Threats and promises are the most common arguments used in real-life negotiations (Kraus et al. 1998). Threats are communications of intent to punish the target if it fails to accept a particular course of action or proposal (Carnevale 1994). Threats are more effective the larger the penalty threatened and the greater their credibility. Credibility refers to the perceived likelihood that a threat will be carried out. There are of number of virtues associated with the use of credible threats, notably they can reduce the attractiveness of the no agreement option for the target and, consequently, can be successful at eliciting concessions from it. However, threats have their down side: they tend to generate resentment and resistance, and they tend to be reciprocated, which can produce counter-threats and destructive conflict spirals (Pruitt and Kim 2004).

Promises are the mirror image of threats—they commit a negotiator to reward a target for compliance to negotiator's demands instead of punishing the target for non-compliance (Pruitt and Carnevale 1993). Credibility is an issue with promises as with threats. Promise credibility is enhanced by past consistency and by past accommodative behaviour. Promises have few of the deficiencies associate with threats. If they work, they tend to build credit rather than resentment. Also, they evoke less resistance than threats and are more flexible in the face of failure (if they do not work, negotiators can simply try another tactic). Promises also have their down side: an effective promise costs a negotiator whatever was promised, the fulfilment of a promise may, paradoxically, make it less likely to work in the future, and the repeated tender and fulfilment of promises may create the problem of undue dependency. Combinations of threats and promises are commonly used in negotiation. Threats are used to emphasize boundaries: to signal the limits below which negotiators refuse to concede and to bolster the restrictions that negotiators wish to place on the other party's behaviour. Promises are used to encourage cooperation within the confines of these boundaries (Pruitt and Kim 2004).

Summary of the state of the art and exemplar negotiation model. Argumentation was originally studied by philosophers and logicians in an attempt to understand the way humans interact with one another (Walton and Krabbe 1995). Recently, the study of argumentation has become of great interest to AI researchers, particularly in defeasible reasoning (Prakken and Vreeswijk 2002), and in multi-agent systems (Reed and Norman 2004). The focus of much of the argumentation research in multi-agent systems was the application of argumentation to negotiation. Sycara (1987, 1990), for example, developed the Persuader system for resolving labor-management conflicts. The system generates an initial compromise and submits it to the parties, who can either accept or reject it. The possible responses to rejection are: (1) to improve the compromise, or (2) to persuade the rejecting party to change its evaluation of the compromise. The latter task involves the construction of persuasive arguments either from scratch or by selecting and adapting previously used arguments through case-based reasoning (see also Sect. 6.5). Kraus et al. (1998) extended and formalized the work of Sycara—they developed a mental model for autonomous agents, proposed an axiomatization system for argumentation, and described a negotiation framework for

creating agents that are able to argue. The mental model uses modal operators for beliefs, desires, intentions, and goals, captures some properties of these modalities, and allows the definition of different types of agents. The axiomatization system formalizes six types of argument, namely threats, promises to future rewards, and appeals to: past rewards, precedents, prevailing practice, and self-interest. The authors also discussed the evaluation of arguments (particularly, threats) and the suitability of arguments for different agent types. Sierra et al. (1998) proposed an approach to construct argumentation-based agents—they extended a protocol for trading proposals with a series of illocutionary moves allowing for the passing of arguments. They proposed three languages for each agent: a base language L (for expressing sentences about negotiation), a meta-language ML (for expressing preferences over formulae in L), and a communication language CL (for communication between agents). The language CL accounts for two sets of illocutionary particles: one set allows the expression of acts about negotiation (notably, offers and counter-offers), and the other set allows the expression of acts about argumentation (particularly, threats, rewards and appeals).

These and other pieces of work on argumentation-based negotiation may be viewed as attempts to marry the exchange of offers with the exchange of arguments. They do this by extending the range of permissible locutions beyond the exchange of offers and by permitting these locutions to transmit arguments for propositions in addition to the propositions themselves. This permits great flexibility since, for instance, it makes possible to persuade agents to change their view of an offer during the course of negotiation. However, little attention was given to a number of important issues, notably:

1. Do the agents use the same protocol for exchanging offers and arguments? If the protocols are different, how do agents know when to move from one protocol to another?
2. How to relate negotiation strategies with arguments? How to define sequences of strategies involving the submission of both offers and arguments (e.g., threats and promises) in discernible patterns?
3. How to communicate arguments? Do the agents communicate arguments as every other piece of information or do arguments constitute separate locutions?
4. Under what conditions does argumentation effectively improve negotiation performance? How to identify effective criteria for selecting the best arguments for specific negotiation settings?

Amgoud et al. (2000a,b) suggested a separate protocol for argumentation which defines how agents generate and respond to arguments based upon what they know. However, the authors pointed out that the protocol "hard-wires" the attitude that the agents take when negotiating with others, defining, for instance when arguments are found to be persuasive, and when their grounds can be questioned. This may lead to agents that are rather inflexible in their argumentation stance (though more flexible than agents which cannot argue). Therefore, as Jennings et al. (2001) pointed out, there is a need to develop more flexible argumentation protocols and to investigate more closely the relation between negotiation and argumentation protocols.

Parsons et al. (1998) and Parsons and McBurney (2003) acknowledged that the approach proposed by Sierra et al. (1998) for constructing argumentation-based agents has a good deal of similarity with the kind of argumentation humans engage in every day. However, they pointed out that the approach does not address the first set of issues raised above and it makes considerable demands on any implementation, mainly because it appeals to very rich representations of the agents and their environments. Therefore, they adopted a different approach—they considered that the agents use argumentation both for internal reasoning

and as a means of relating what they believe and what they communicate. The agents are only permitted to make assertions and to accept assertions that are backed by arguments, i.e., the formation of arguments becomes a pre-condition of the locutions of the communication language and the locutions are linked to the agents' knowledge bases (but see below).

Furthermore, Jennings et al. (2001) observed that argumentation-based models add considerable overheads to the negotiation process, not least in the construction and evaluation of arguments. Accordingly, they pointed out that agents which can argue in support of their negotiations will only ever represent a small, though important, class of automated negotiators. Karunatillake and Jennings (2005) also pointed out that no real attention was given to the real impact of the decision made by the agents to resolve their conflicts by arguing. Against this background, the remainder of this subsection presents a detailed description of a prominent argumentation-based model proposed in the literature.

Parsons and Jennings (1996) and Parsons et al. (1998) proposed a framework for autonomous agents that accounts for individual behavior (a mental model relating perception to action) and social behavior (argumentation-based negotiation). They claimed the following contributions towards the goal of developing autonomous agents with negotiation competence:

1. a multi-context approach to build autonomous agents;
2. a generic framework for negotiation;
3. a logic-based system of argumentation;
4. a framework for argumentation-based negotiation.

The multi-context approach is founded upon the natural correspondence between multi-context systems (Giunchiglia and Serafini 1994) and the modularity of agent architectures. The main components of multi-context agents are: (1) units (structural entities), (2) logics (declarative languages), (3) theories (sets of formulae written in the logic associated with a unit), and (4) bridge rules (rules of inference which relate formulae in different units). The negotiation framework characterizes negotiation as a dialogue involving the exchange of proposals, critiques, explanations (justifications that the agents supply for their positions), and meta-information (information to focus the local search by agents for solutions). The argumentation system is based upon that proposed by Fox et al. (1992) and Krause et al. (1995) and works by constructing series of logical steps (arguments) for and against propositions of interest. The system formalizes the notion of argument, defines some relationships between arguments (notably, rebut and undercut), and proposes a hierarchy of acceptability classes for arguments. The framework for argumentation-based negotiation characterizes negotiation as a process that involves essentially the construction, evaluation, and submission of arguments and, as a result, falls within the generic framework just described. The authors claimed that the multi-context approach supports the development of agent architectures which have a clear link between their specification and their implementation (see Sabater et al. 2002, for a detailed discussion). Also, they claimed that the multi-context approach and the argumentation system constitute the basic support to build autonomous agents capable of reasoning and negotiating by arguing (see also Table 2, in "Appendix").

6.3 Learning in negotiation

Negotiation is a learnable process—most negotiators can gain greater confidence in their ability and improve their skills with a few lessons, a bit of coaching, some tips on how to negotiate better, and repeated experience (Lewicki et al. 1999). Learning is the acquisition of new knowledge and motor and cognitive skills and the incorporation of the acquired

knowledge and skills in future agents' activities, provided that this acquisition and incorporation is conducted by the agents themselves and leads to an improvement in their performance (Sen and Weiss 1999).

The following four learning methods have been commonly used by negotiators (Nadler et al. 2003; Thompson 2005):

1. principle-based or didactic learning;
2. learning by feedback or via information revelation;
3. learning by analogy or analogical learning;
4. observational learning or imitation.

Principle-based or didactic learning involves mainly providing negotiators with general principles of negotiation, and then asking them to solve problems that require the use of the same principles. Learning by feedback or via information revelation is based on the premise that negotiators can learn through their own behaviour—simply put, past or current outcomes may influence subsequent outcomes. Learning by analogy or analogical learning involves noticing that solutions to (source) problems from the past are relevant, and then mapping the elements from those solutions to produce solutions to (current or target) problems. Observational learning or imitation is based on the premise that negotiators can improve their own skills by observing those of others.

Several studies suggest that learning plays a key role in negotiation. Thompson and DeHarpport (1994), for example, found that revealing information about the opponents' priorities and preferences at the conclusion of a completed negotiation resulted in superior transfer, or performance on a subsequent negotiation task, compared to providing outcome information alone. Ross and Kilbane (1997) revealed that a close connection between a principle and relevant examples is crucial for learners taking advantage of abstract principles. Loewenstein et al. (1999) and Thompson et al. (2000) demonstrated that analogical comparison of solutions of past problems and subsequent derivation of abstract schemas on the basis of their commonalities encouraged recognition of target problems' structural features, leading to advantages in negotiation performance. Nadler et al. (2003) examined the efficacy of the four learning methods just outlined and found that observational and analogical learning led to superior joint outcomes, compared to a baseline condition of learning through experience alone. Information revelation and didactic learning were not significantly different from any other condition.

Summary of the state of the art and exemplar negotiation model. AI researchers have traditionally adopted a "divide and conquer" strategy for developing autonomous agents and, as a result, they have focused on isolated cognitive capabilities. Automated negotiation and machine learning were originally studied separately. In the last several years, however, the mutual ignorance of machine learning and automated negotiation has started to disappear. Sycara (1987, 1993), for example, developed the Persuader system which uses case-based reasoning to learn from its experience (see also Sect. 6.5). Prasad et al. (1998) presented a multi-agent parametric design system called L-TEAM where a set of agents learn their organizational roles in negotiated search for mutually acceptable designs. L-TEAM produced better results than its non-learning predecessor TEAM (Lander 1994). Oliver (1997), Matos et al. (1998), Gerding et al (2003) and Gerding and Poutré (2006) demonstrated that artificial adaptive agents using evolutionary algorithms can learn effective strategies for different negotiation situations.

Coehoorn and Jennings (2004) extended the negotiation model developed by Faratin (2000) and Faratin et al. (2002) by incorporating a new method (kernel density estimation) to

learn the preferences of an opponent in bilateral negotiation. The extended model led to better performance (in terms of payoff and time of agreement) than the original model. A number of researchers presented models that incorporate Bayesian learning mechanisms (e.g., Zeng and Sycara 1998; Nguyen and Jennings 2004). Coehoorn and Jennings (2004) pointed out, however, that the agents often need a priori information to determine the probability of a set of hypotheses about their opponents' preferences, which is difficult to obtain due to its private nature.

Despite these and other relevant efforts to develop autonomous negotiating agents that are able to learn, there are a number of questions still waiting to be addressed more thoroughly. We highlight the following:

1. How to combine different learning methods (e.g., learning from experience with other learning methods)?
2. Under what conditions does learning effectively improve negotiation performance?
3. Which are the most effective methods of learning? Likewise, how to assess the effectiveness of different learning methods?

Probably, the most widely used learning paradigm is Bayesian learning and we present next a representative example of this line of work.

Zeng and Sycara (1998) presented a sequential decision-making model of negotiation, called Bazaar, where the agents are self-interested, computationally bounded, and able to learn. The main components of the model are:

1. a negotiation protocol;
2. a set of negotiation strategies;
3. a learning mechanism.

The negotiation protocol is the alternating offers protocol (Osborne and Rubinstein 1990). The negotiation strategies define sequences of actions (proposals) that end with either the accept action or the quit action. The authors emphasized the learning aspects of Bazaar—the agents keep histories of their interactions and update their beliefs, using a Bayesian updating mechanism, after observing the environment and the behaviour of the other negotiating agents. Specifically, prior to the negotiation, each agent has a probability distribution defined over a set of beliefs about both the environment and the other negotiating agents (e.g., beliefs about the reservation prices or payoff functions). During the give-and-take of negotiation, in particular after receiving a proposal, each agent updates the probability distribution before entering the next stage of negotiation (the reply to the proposal is based on newer information). The updating also occurs when each agent discovers new externally available information (e.g., if a buyer finds out that the overall supply of a particular good under negotiation is experiencing a tremendous increase, its estimated of a supplier's reservation price might drop without even receiving any new offers from the supplier). The authors claimed that Bazaar handles multi-party and multi-issue negotiation, supports an open-world model, and addresses the issue of learning during negotiation (reassessment of perceptions about both the environment and the opponents). They presented both theoretical analysis and initial experimental results showing the benefits of learning in simple negotiation situations. The main features of this work are summarized in Table 2 (see "Appendix").

6.4 Selection of new strategies

Negotiation strategies refer to the patterns or plans to accomplish the goals of the parties (Lewicki et al. 2003). Effective negotiators often make a conscious analysis of both the

negotiation situation and the other parties and actively prepare initial strategies that match this judgment. Also, effective negotiators often update their judgement as negotiation unfolds—information received during negotiation frequently causes them to change their perception of the structure of the negotiation situation and their opinion about the other parties. Hence, effective negotiators may move back and forth among different strategies in discernible patterns (Putnam 1990).

Now, as stated in Sect. 5.3, there are three fundamental groups of strategies commonly used by negotiators: contending, problem solving, and concession making. The strategies in each group are pure strategies, typically at odds with the mixture of elements that characterize the majority of negotiation situations. The number of different sequences that theoretically can be defined with the strategies of these three groups is enormous. The discussion and analysis of the most important sequences is, therefore, beyond the scope of this paper. However, for the sake of readability, and also to illustrate some problems, we discuss next a sequence of strategies that has attracted much attention in negotiation research: moving through stages.

Some real-life negotiations follow a two-stage sequence—the first stage involves a combination of contending and concession making, while the second stage involves heavy problem solving (Pruitt and Carnevale 1993). Negotiators typically start with large demands backed up by contentious tactics (e.g., positional commitments, persuasive maneuvers, and attacking arguments). They make efforts to demonstrate firmness around their key proposals and to persuade the opponent(s) to move toward them. They also make a few easy concessions to show good will and thus keep the negotiation going. The transition between the two stages often involves a sense of hurting stalemate—the parties are not willing to make further concessions and the continued use of contentious tactics puts pressure on resources and runs the risk of failure to reach agreement. Out of this stalemate often comes an effort to work together in search for a mutually acceptable agreement. Negotiators make efforts to separate invention of solutions from decisions through brainstorming, exchanging proposals, sharing information, and deciding on the terms of agreement.

Hard evidence of this two-stage sequence was observed in several studies of labour-management negotiation (Morley and Stephenson 1977). Related sequences of stages have also been observed in international negotiation (Snyder and Diesing 1977) and community mediation (McGillicuddy et al. 1987). Some researchers postulate an additional, competitive, stage beyond these two—the negotiation begins competitively as parties seek to identify and differentiate issues in dispute, becomes cooperative as parties enter into problem solving, and end competitively as parties vie for advantage in the final agreement. However, the support for reviving competitive activity in the final stage is mixed. Although some researchers observe that competitive activity resurges in the final stage when negotiators affirm their gains, other researchers report a continual increase of problem solving activity to facilitate reaching a final agreement (Putnam 1990).

Summary of the state of the art and exemplar negotiation model. AI researchers have traditionally concentrated on the formalization and evaluation of negotiation strategies. In particular, as pointed out in Sect. 5.3, AI researchers have developed several models that include libraries of negotiation strategies (notably, concession and problem solving strategies) and have investigated the effectiveness of these strategies in various negotiation situations. However, most researchers have neglected the important pre-negotiation step of selecting appropriate strategies for specific negotiation situations. Also, most researchers have treated strategies as rigid or static elements of negotiation, i.e., elements that do not change during negotiation. There have been few attempts to develop models that incorporate effective approaches to dynamic strategic choice—most models do not support the selection

of new strategies as negotiation unfolds. The work of Nguyen and Jennings (2003, 2004, 2005) is a significant exception (recall the exemplar model of Sect. 5.3 and see also Sect. 7.1, below). The remainder of this section presents another piece of work that allows agents to move back and forth between a mechanism that defines a range of concession strategies and an algorithm that enables the generation of trade-offs (i.e., concessions on some issues and higher demands on other issues).

Faratin et al. (1998), Faratin (2000) and Faratin et al. (2002) presented a model for bilateral, multi-issue service-oriented negotiation. The authors considered a negotiation situation where an agent (a client) requires a service to be performed on its behalf by some other agent (a server). The agents negotiate the terms and conditions under which the service will be executed (e.g., the price of the service and the time at which it is required). The main components of the model are:

1. a bilateral negotiation protocol;
2. a responsive mechanism;
3. a trade-off mechanism;
4. an issue manipulation protocol;
5. an issue set manipulation mechanism;
6. a meta-strategy mechanism.

The bilateral negotiation protocol is a natural extension of the contract net protocol (Smith 1980) permitting an iterative exchange of offers and counter-offers. The responsive mechanism defines a range of concession strategies and tactics that every agent can employ to generate offers and counter-offers. Strategies are functions that determine the linear combinations of tactics to be used throughout negotiation. Tactics are functions that determine new values for each issue at stake in negotiation. Three groups of tactics were presented: (1) time dependent (functions of time), (2) resource dependent (functions of limited resources), and (3) behavior dependent or imitative (functions of the opponent's behavior). The trade-off mechanism defines a linear algorithm enabling every agent to make offers that have the same score as previous offers, but are more acceptable to the opponent (an agent demands more on some issues and simultaneously lowers its acceptance level on other issues). The algorithm performs an iterated hill-climbing search to explore the space of possible trade-offs and uses the notion of fuzzy similarity (Zadeh 1971) to approximate the preferences of the opponent. The issue manipulation protocol is similar to the negotiation protocol and specifies the possible tasks that every agent can perform to dynamically include or retract issues from the set of negotiation issues. The issue set manipulation mechanism defines two operators: add and remove. These operators assist an agent in selecting an issue to include or retract from the current set of issues, respectively. The meta-strategy mechanism is mainly responsible for deciding which of the other three mechanisms should be employed throughout negotiation. To this end, the authors discussed the role and rationale of meta-strategies that combine different mechanisms towards an outcome (e.g., a meta-strategy that continuously switches between the responsive mechanism and the trade-off mechanism). However, they postponed the formal treatment of meta-strategies to future work. Table 2 summarizes the key features of this piece of work (see "Appendix").

6.5 Impasse resolution

Negotiation may end in a settlement, wherein the parties mutually agree to a proposal, or in an impasse, a deadlock, a stalemate, or a breakdown, wherein the parties do not reach a settlement (Thompson 2005). An impasse is a condition or state of negotiation in which there

is no apparent quick or easy resolution—the parties are unable to create mutually advantageous deals that satisfy their aspirations (Lewicki et al. 2003). Productive dialogue stops. The parties may continue interacting, but the communication is usually characterized by trying to force individual positions, discussing the unreasonable positions and uncooperative behaviour of the others, or both. Issues are viewed is such a way that the parties do not believe that there is any possible compatibility between them, or they cannot find a middle ground where agreement is possible.

Through any number of different avenues—polarization of positions and refusal to compromise, the issuance of ultimatums, or breakdowns in communication—negotiation frequently hits an impasse. The parties should take jointly several actions to move negotiation back to a productive pace. Lewicki et al. (2003) suggested the following procedure for resolving highly polarized impasses:

1. synchronize the de-escalation of hostility and improve the accuracy of communication;
2. keep the number of issues under control so that issues are managed effectively, large issues are divided into smaller ones, and new issues are not careless added;
3. establish a common ground (e.g., by determining a common framework for approaching the negotiation problem, including managing time and deadlines);
4. make proposals more attractive for joint resolution (e.g., by understanding the needs of the opposing negotiators and devising proposals that will meet those needs).

The authors observed that these actions represent self-help for negotiators—they are possible ways that the parties themselves can use to improve the odds that successful resolution can occur and to overcome intractability. However, as they pointed out, frequently the parties cannot execute some or all of the actions effectively by themselves and the intervention of third parties—that is, individuals or agencies other than the disputants themselves—may become necessary. Third-party intervention should be avoided as long as progress is occurring or is likely to occur within reasonable limits of time and other resources. However, third-party intervention may be a productive way (if not the only way) to get difficult negotiations back on track.

Mediation is probably the most common type of third-party intervention. Mediators have the ability to meet with the parties individually, secure an understanding of the issues in dispute, identify areas of potential compromise, and encourage the disputants to make concessions toward agreement. However, they have no formal power over the outcome, and they cannot resolve a dispute on their own or impose a solution—they let the parties themselves develop and endorse a final agreement. Mediation, like negotiation, often proceeds through distinct phases or stages (see e.g., Lovenheim 1989; Moore 1996). Most phase models fit into a general structure of four groups of stages (Lewicki et al. 2003): premeditation preparation, beginning stages, middle stages, and ending stages. In the premeditation stages, mediators are most concerned with understanding the nature of the dispute and with securing acceptance by the parties. The beginning stages are characterized by establishing ground rules and setting an agenda. As mediation progresses, mediators start to manage the exchange of proposals and counterproposals, inventing solutions, and pressing the parties to make concessions. In the ending stages, mediators will bring the parties together to endorse a final agreement or to publicly announce their settlement.

Mediation is effective in general—disputants are frequently satisfied, agreement is commonly reached, and compliance rates are high (Pruitt and Kim 2004). Research has pinpointed certain characteristics of mediation situations and certain kinds of mediator behaviors that are especially likely to produce success. Various studies have shown that mediation is more effective when the parties are committed to mediation, conflict is moderate rather than intense,

the parties are relatively equal in power, the issues do not involve a general principle, and the parties are highly motivated to reach settlement (Kressel and Pruitt 1989; Carnevale and Pruitt 1992). Mediator behaviors positively related to successful mediation include creating and controlling the agenda, helping the parties establish priorities among the issues, and maintaining a firm control over the mediation process (Zubek et al. 1992; Deutsch 1994).

Summary of the state of the art and exemplar negotiation model. Traditionally, AI researchers have focused on understanding and formalizing "successful" negotiations—most researchers have assumed that negotiations result in agreement, i.e., negotiations do not become difficult to resolve and contentious to the point of impasse. There are, however, some researchers that have studied various techniques that autonomous agents can use to resolve impasses on their own. Lopes et al. (2002), for example, presented a model that formalizes a structure for the problem under negotiation and supports the dynamic change of that structure to achieve movement towards agreement (problem restructuring). The authors pointed out that problem restructuring allows the addition and removal of issues during the course of negotiation, thus facilitating the resolution of impasses and increasing the parties' willingness to settlement. However, they observed that problem restructuring is a highly creative and challenging task and postponed its formal treatment to future work (but recall the exemplar model of Sect. 5.1). Faratin (2000) and Faratin et al. (2002) presented a model that incorporates a (partially developed) mechanism to assist agents in dynamically including or retracting issues from the set of negotiation issues. The authors acknowledged the usefulness of the mechanism to escape negotiation deadlocks. However, they pointed out that the mechanism is complex and deferred to future work both the specification of an issue manipulation algorithm and its empirical analysis (recall the exemplar model of the previous subsection).

There are also some researchers that have studied third-party approaches to resolve impasses, notably mediated negotiation (see e.g., Gordon and Karacapilidis 1997; Klein et al. 2003). However, Rahwan et al. (2004) pointed out that mediated negotiation is a significantly unexplored topic and noted the absence of a grounded theory of mediation in multi-agent negotiation.

At present, despite these and other relevant pieces of work, the study of "difficult" negotiations and the implementation of approaches to impasse avoidance and/or resolution are still very much in its infancy. We highlight the following issues:

1. Why and how negotiations become "difficult" to resolve and reach impasse? What are the causes of impasse (e.g., the characteristics of the parties, the nature of the conflict, or the characteristics of the issues in dispute)?
2. How to manage "difficult" negotiations? Which individual approaches are effective? Likewise, which are the specific actions that the agents can take jointly to effectively return negotiations to a productive pace?
3. When is third-party involvement appropriate? Which type of intervention is appropriate?

The remainder of this subsection describes work that incorporates a mediator and supports problem restructuring.

Sycara (1987, 1990, 1991, 1993) developed the Persuader system that incorporates three agents: a mediator, a company and its union. The mediator engages in negotiation with the parties when their goals are in conflict—its input is a dispute, i.e., a situation in which the goals of the parties are incompatible or there is a limitation in the resources needed to fulfil all the goals. The interaction process involves three main tasks:

1. generation of proposals;
2. persuasive argumentation;
3. generation of counter-proposals (modification of rejected proposals).

The mediator generates an initial compromise and submits it to the company and the union. The compromise can be generated by employing either: (1) case-based reasoning, (2) preference analysis, or (3) situation assessment. Case-based reasoning is the central method and involves mainly the retrieval of appropriate precedent cases from a memory of past negotiation experiences, the construction of an initial solution, and the modification of the solution to fit the current problem. Preference analysis consists basically in using the individual utility functions of the agents and selecting the proposal that maximizes the joint payoff and minimizes the payoff difference (the author claimed that this criterion balances maximal gains with equity). This method is used when no appropriate precedent cases are available. Situation assessment consists mainly in accessing domain independent knowledge structures that incorporate general strategies for dealing with exceptional cases. The company and the union receive the compromise and can either accept or reject it. The possible responses to a rejection by one of the parties are: (1) to persuade the rejecting party to agree to the compromise, or (2) to modify the compromise. Persuasive argumentation is the preferred method and involves the construction of arguments either from scratch or by selecting and adapting previously used arguments through case-based reasoning. Compromise modification involves basically the following tasks: (1) asking a rejecting party for feedback information (concerning the objectionable issues, their importance, and the reason for rejection), (2) using the feedback information to select a similar impasse and access the associated modification, and (3) adapting the modification to the current impasse. Problem restructuring may take place during the modification of a compromise. Table 2 summarizes the key features of this piece of work (see "Appendix").

7 Settlement (closing and implementing the agreement)

7.1 Final agreement

The final periods of negotiation are frequently marked by the exchange of specific, substantive proposals, accompanied by concession making when parties converge on points of agreement. The parties often demand a gesture of good faith commitment to the agreement (close the deal) and determine who needs to do what once the documents are signed (implement the agreement). Typically, they are satisfied or happen with the negotiation outcome and, in many instances, proudly describe it. However, closer inspection usually reveals that money was squandered, resources wasted, and potential joint gain untapped (Thompson 2005). Also, not uncommonly parties discover that the agreement is flawed, key points were missed, or the situation has changed and new questions exist. As a result, the deal may have to be reopened, or issues settled by arbitrators or the courts (Greenhalgh 2001).

The negotiation literature commonly identifies three different settlement types: a compromise agreement, an integrative agreement, and a Pareto optimal agreement. As stated in Sect. 5.3, a compromise agreement is achieved when the parties concede to a middle ground on some obvious dimension or dimensions that link their initial proposals. Also, an integrative agreement is an agreement that reconciles (i.e., integrates) the parties' interests and thus provides higher joint benefit than a compromise agreement.

A Pareto optimal agreement is an agreement that lies along the Pareto optimal or efficient frontier, i.e., the locus of achievable joint evaluations from which no joint gains are possible (Raiffa 1982). This means that a Pareto optimal agreement produces the highest joint benefit of the three types of settlement. Unfortunately, negotiators frequently view conflict-laden situations with a fundamentally more distrustful, win-lose attitude than is necessary or desirable and fail to reach integrative or even Pareto optimal agreements. They overlook opportunities for mutually beneficial agreements and instead, settle for outcomes that are worse for them than other available solutions.

Raiffa (1985) presented a promising approach to minimizing this limitation on rationality in negotiation. The basic idea is that negotiators who reach an agreement can ask a third-party to help them find a mutually superior solution (one that is better for all parties). The intervenor objectively analyzes the existence of mutually superior solutions based on its expertise and the information available to the negotiators. In the end, each party would reserve the right to veto the post-settlement settlement proposed by the third-party, and revert to the original agreement. Bazerman et al. (1987) extended Raiffa's approach by arguing that negotiators should look for a post-settlement settlement without the help of a third-party. They observed that skilled negotiators can share information and make the appropriate analytical insights. Also, they claimed that the concept of post-settlement settlement is most useful if thought of in terms of a process instead of an outcome—a process by which the parties reaching an initial agreement have the commitment to continue to search for a superior solution as more information is obtained and analyzed. Ehtamo et al. (1999a, 2001) presented the method of improving directions which formalizes the single negotiation text procedure introduced by Roger Fisher for the Camp David negotiations (see Fisher and Ury 1981, and also Raiffa 1982). Ehtamo et al. (1999b) presented the constraint proposal method for generating Pareto optimal solutions in two-party negotiations and Heiskanen et al. (2001) extended it to multi-party negotiations. Thompson (2005) proposed a strategic framework for reaching integrative and Pareto optimal agreements. The framework has modules for: (1) resource assessment, (2) assessment of differences, (3) construction of offers and tradeoffs, and (iv) renegotiation.

Summary of the state of the art and exemplar negotiation model. AI researchers have traditionally concentrated on the study of agreements that involve different levels of commitment and different sanctions for decommitment. Most researchers have developed models that impose unbreakable commitments, i.e., models that do not incorporate any form of decommitting possibility. The commitments cannot be broken, no matter how future events unravel. Excelente-Toledo et al. (2001) argued, however, that the imposition of unbreakable commitments can lead to inefficient behavior. Sandholm and Lesser (2001) also pointed out that contracts can be profitable when viewed ex ante, but not necessarily when viewed ex post. They argued that full-commitment contracts are unable to capitalize on the potential gains that future events can provide. Accordingly, a number of researchers have acknowledged the potential provided by contingency contracts (Raiffa 1982) and developed models that allow commitments to be dropped to accommodate future events. The commitments are made contingent on specific future events. However, Sandholm and Lesser (2001, 2002) pointed out that there are at least three major problems in using contingency contracts: (1) agents often do not know all possible future events beforehand and, therefore, they cannot always use contingency contracts optimally, (2) contingency contracts are useful in anticipating a small number of key events, but become cumbersome as the number of future events to monitor increases, and (3) some future events may be observable by only one of the contract parties (that may lie if the events are associated with disadvantageous contingencies). Therefore, they proposed leveled-commitment contracts that allow unilateral decommitting—any agent can decommit

from a contract by simply paying a fixed decommitment penalty. Excelente-Toledo et al. (2001) and Nguyen and Jennings (2005) argued, however, that leveled-commitment contracts do not explicitly take into account the ongoing cost of participating in the negotiation process. Accordingly, they introduced the notion of variable penalty contracts, as an extension to leveled-commitment contracts, that take into account the stage of the negotiation process at which a commitment is broken.

However, despite these and other relevant efforts to formulate agreements that accommodate future events, AI researchers have paid little attention to a number of issues related to the analysis and improvement of final agreements. We highlight the following:

1. When to reopen a deal and try a second-agreement negotiation (renegotiation)?
2. How to search for post-settlement settlements (without the help of a third party)?
3. How to integrate the process of moving toward agreement (actual negotiation) with the process of searching for mutually superior solutions (renegotiation)? Likewise, how to view actual negotiation and renegotiation as a single process that culminates in a mutually superior solution?

There have been few attempts to develop models that incorporate effective renegotiation approaches. Against this background, the remainder of this section presents a piece of work that allows agents to search for superior outcomes.

Nguyen and Jennings (2003, 2004, 2005) presented a model that handles one-to-many negotiation in service-oriented contexts—a buyer engages in multiple concurrent bilateral negotiations with a set of sellers capable of providing a specific service (recall the exemplar model of Sect. 5.3). The main components of the buyer are a coordinator, a number of negotiation threads (one per seller), and a commitment manager. The coordinator is mainly responsible for coordinating all the negotiation threads, choosing appropriate strategies for the threads, and dynamically changing their strategies. The negotiation threads deal directly with the sellers and are mainly responsible for making offers and for deciding whether or not to accept the offers made by them. The commitment manager is responsible for handling any issue that is related to commitment and decommitment—it is invoked when a thread needs to decide whether or not to accept an offer or when a seller decides to renege from a committed deal. Specifically, the buyer could make a number of intermediate deals with various sellers (where each such deal is a temporary agreement with a particular seller). However, during the course of negotiation, if either the buyer or any seller decides to break a committed deal (for whatever reason), it has to pay a decommitment fee to its opponent. The fee is dynamically calculated as a percentage of the utility of the deal and is also based on the time when the contract is broken. Thus, when presented with a potential agreement from a specific seller, the buyer has to decide whether it should accept or reject it. To this end, the buyer takes into account the following:

1. if the buyer already has a commitment with another seller, the utility gained by taking the new offer must be greater than that of the current deal, after having paid the decommitment fee;
2. the degree of acceptance for the new offer must be over a predefined threshold; the threshold specifies how the buyer accepts offers, in particular whether it is greedy (tends to accept any possible deal) or patient (only deals that provide a certain expected utility value are accepted).

The buyer can commit to more than one deal at any time and, later on, select the deal that has the highest utility value and decommit from the others. This avoids the risk associated with

committing to only one deal which can then be revoked near the deadline, leaving the buyer with insufficient time to find a replacement. The main features of this work are summarized in Table 2, in "Appendix".

8 Conclusion

This article has pointed out that automated negotiation systems are becoming increasingly important and pervasive and has argued that, at present, there is a need to build a generic framework to define and characterize the essential features that are necessary to conduct automated negotiation. The article has also claimed that the development of such a framework can be an important step to identify the core elements of autonomous negotiating agents, to provide a coherent set of concepts related to automated negotiation, to understand the interrelationships of disparate research efforts, to assess progress in the field, and to highlight new research directions, constituting an important tool for the development of more sophisticated negotiating agents.

Accordingly, the article has introduced a generic framework for automated negotiation. It has described, in detail, the components of the framework (with reference to work from artificial intelligence and various fields of the social sciences), has presented an appraisal of the relative merits and drawbacks of the majority of negotiation models proposed in the literature (according to each component of the framework), and has described a number of prominent models. It has also raised a number of new questions and highlighted some of the major challenges for future automated negotiation research. To recap:

1. AI researchers have traditionally not recognized and explored the role conflict plays in driving agent interaction and negotiation. Few researchers have acknowledged the inevitability of conflict in multi-agent systems, developed methods for conflict detection, and explored the presence of conflict to drive negotiation;
2. AI researchers have focused both on bilateral negotiation and one-to-many negotiation. Aside from the formal work on bargaining, auctions, market-oriented programming, contracting and coalition formation, little work has addressed either negotiation involving more than two parties attempting to achieve a collective consensus or negotiation between teams or groups;
3. Pre-negotiation research has been conducted in a slightly fragmented manner—rather than approaching the content of pre-negotiation plans as a coherent structure, AI researchers have focused in a few activities in isolation from other activities (notably, the formal representation of the issues to be deliberated, the formulation of their relative importance, and the definition of targets and limits). Little work has addressed the formalization of the problem under negotiation, the identification of the issues, the definition of the negotiating agenda, the process of acquiring information about the opposing negotiators, and the selection of appropriate initial strategies for specific negotiation situations;
4. AI researchers have traditionally concentrated on modeling the offer/counter-offer process. Recently, an increasing number of researchers have also started to pay attention to the role of feedback information in negotiation (particularly in guiding the generation of offers that are more likely to be acceptable) and the role arguments play in negotiation (notably, persuasive arguments, threats, and promises). However, little work has addressed the communication of key pieces of information that promote the development of mutually superior solutions, the relation between negotiation and argumentation proto-

cols, the definition of sequences of strategies involving the submission of both offers and arguments in discernible patterns, and the identification of the conditions under which the communication of feedback information and/or specific arguments effectively improves negotiation performance;

5. Learning in negotiation is still a young area of research—the mutual ignorance of automated negotiation and machine learning has disappeared, but few researchers have attempted to develop negotiation models that formalize important learning methods studied in the social sciences and commonly used by human negotiators. Also, little work has attempted to combine different learning methods and determine the most effective methods in various negotiations situations;
6. AI researchers have treated strategies as rigid or static components of negotiation—they have considered strategies as elements that do not change during negotiation, rather than considering them as elements that can and typically do change during negotiation. There have been few attempts to develop models that incorporate effective approaches to dynamic strategic choice;
7. AI researchers have traditionally focused on modeling "successful" negotiations—most researchers have assumed that negotiations result in agreement, i.e., negotiations do not become contentious to the point of impasse. Few researchers have attempted to understand the nature of negotiations that are "difficult" to resolve, formalize the causes of impasses, and implement effective approaches to impasse resolution;
8. Renegotiation has traditionally been oversimplified—most researchers have assumed that agreements involve unbreakable commitments. Recently, a number of researchers have started to study agreements that accommodate future events (i.e., agreements that involve different levels of commitment and different sanctions for decommitment). However, there have been few attempts to develop models that incorporate effective renegotiation approaches—little work has addressed the use of the first agreement as a plateau that can be improved upon in a follow-up negotiation effort.

In identifying these theoretical and practical issues, we want to spark interest in areas of research that we believe will help to fill gaps in our current knowledge of automated negotiation. We do not intend to dictate what research should be done, but we do want to point to some research areas that seem quite underdeveloped and that offer exciting opportunities for future study. Finally, we regard research on automated negotiation as an exciting, vigorous tradition that has generated many useful ideas and concepts (leading to important theories and systems). Nevertheless, this field is still immature. There is much further work to be done, as we have tried to indicate in this paper.

Acknowledgments This work has been partially supported by the Technology and Science Foundation (FCT) grant SFRH/BPD/23921. The authors also wish to acknowledge the valuable comments and suggestions made by two anonymous reviewers.

Appendix: Analysis of prominent negotiation models

See Table 2.

Table 2 Analysis of a representative sample of the most prominent negotiation models

		Katia Sycara	Lander and Lesser	Zeng and Sycara	Parsons et al.	Faratin et al.	Lopes et al.	Nguyen and Jennings	Li et al.
Preliminaries	1. Negotiating parties (number of parties)	3	*n*	*n*	*n*	2	*n*	*n*	*n*
	2. Social conflict (detection and exploration)	◉	■	□	□	□	■	□	□
Pre-negotiation	3. Structuring of personal information (definition and execution of key tasks)	□	□	□	□	□	◉	□	□
	4. Analysis of the opponents (gathering and use of key information)	□	□	◉	□	□	□	◉	◉
	5. Definition of the protocol and selection of the initial strategy	□	□	□	□	□	□	■	□
Actual negotiation	6. Exchange of offers and Feedback Information	■	■	□	■	□	◉	□	□
	7. Argumentation (exchange of threats, promises, etc.)	■	□	□	■	□	□	□	□
	8. Learning (in negotiation)	■	□	■	□	□	□	◉	◉
	9. Dynamic strategic choice (selection of new strategies)	□	□	□	□	◉	□	◉	□
	10. Impasse resolution	■	□	□	□	◉	◉	□	□
Renegotiation	11. Analysis and improvement of the final agreement	□	□	□	□	□	□	■	□

The models either address (■), partially address (◉), or do not address (□) every issue

References

Amgoud L, Maudet N, Parsons S (2000a) Modeling dialogues using argumentation. In: 4th international conference on multi-agent systems. IEEE Computer Society Press, pp 31–38

Amgoud L, Parsons S, Maudet N (2000b) Arguments, dialogue, and negotiation. In: 14th European conference on artificial intelligence. IOS Press, pp 338–342

Bartolini C, Preist C, Jennings N (2005) A software framework for automated negotiation. In: Choren R et al. (eds) Software engineering for multi-agent systems III: research issues and practical applications. Springer, Heidelberg, pp 213–235 (LNCS 3390)

Bazerman M, Neale M (1992) Negotiating rationality. Free Press, New York

Bazerman M, Russ L, Yakura E (1987) Post-settlement settlements in two-party negotiations. Negot J 3: 283–292

Bazerman M, Mannix E, Thompson L (1988) Groups as mixed-motive negotiations. In: Lawler E, Markovsky B (eds) Advances in group processes, vol 5. JAI Press, London, pp 195–216

Brett J (1991) Negotiating group decision. Negot J 7:291–310

Bussmann S, Jennings N, Wooldridge M (2004) Multiagent systems for manufacturing control. Springer, Heidelberg

Butler J (1999) Trust expectations, information sharing, climate of trust, and negotiation effectiveness and efficiency. Group Organ Manag 24:217–238

Carnevale P (1994) Negotiation. In: Ramachandran V (ed) Encyclopedia of human behavior, vol 3. Academic Press, New York, pp 271–281

Carnevale P, Pruitt D (1992) Negotiation and mediation. In: Rosenzweig M, Porter L (eds) Annual review of psychology, vol 43. Annual Reviews Inc., Palo Alto, pp 531–581

Castelfranchi C (2000) Conflict ontology. In: Müller H, Dieng R (eds) Computational conflicts. Springer, Heidelberg, pp 21–40

Chaiken S, Gruenfeld D, Judd C (2000) Persuasion in negotiations and conflict situations. In: The handbook of conflict resolution: theory and practice. Jossey-Bass, San Francisco, pp 144–165

Coehoorn R, Jennings N (2004) Learning an opponent's preferences to make effective multi-issue negotiation trade-offs. In: 6th International conference on e-commerce. ACM Press, New York, pp 59–68

Deutsch M (1994) Constructive conflict resolution: principles, training and research. J Soc Issues 50:13–32

De Dreu C, Weingart L, Kwon S (2000) Influence of social motives on integrative negotiation: a meta-analytic view and test of two theories. J Pers Soc Psychol 78:889–905

Ehtamo H, Verkama M, Hämäläinen R (1999) How to select fair improving directions in a negotiation model over continuous issues. IEEE Trans Syst Man Cybern C Appl Rev 29:26–33

Ehtamo H, Hämäläinen R, Heiskanen P, Teich J, Verkama M, Zionts S (1999) Generating Pareto solutions in a two-party setting: constraint proposal methods. Manag Sci 45:1697–1709

Ehtamo H, Kettunen E, Hämäläinen R (2001) Searching for joint gains in multi-party negotiations. Eur J Oper Res 130:54–69

Excelente-Toledo C, Bourne R, Jennings N (2001) Reasoning about commitments and penalties for coordination between autonomous agents. In: 5th international conference on autonomous agents. ACM Press, New York, pp 131–138

Faratin P (2000) Automated service negotiation between autonomous computational agents. Dissertation, Queen Mary & Westfield College, London

Faratin P, Sierra C, Jennings N (1998) Negotiation decision functions for autonomous agents. J Robot Auton Syst 24:159–182

Faratin P, Sierra C, Jennings N (2002) Using similarity criteria to make issue trade-offs in automated negotiations. Artif Intell 142:205–237

Fasli M (2007) Agent technology for e-commerce. Wiley, Chichester

Fatima S, Wooldridge M, Jennings N (2005) A comparative study of game theoretic and evolutionary models of bargaining for software agents. Artif Intell Rev 23:185–203

Fisher R, Ury W (1981) Getting to yes: negotiating agreement without giving in. Arrow Books Limited, London

Fisher R, Ertel D (1995) Getting ready to negotiate: the getting to yes workbook. Penguin Books, London

Fox J, Krause P, Ambler S (1992) Arguments, contradictions, and practical reasoning. In: 10th european conference on artificial intelligence. Wiley, Chichester, pp 623–627

Fox M, Barbuceanu M, Teigen R (2000) Agent-oriented supply chain management. Int J Flex Manuf Syst 12:165–188

Garrido L, Sycara K (1996) Multi-agent meeting scheduling: preliminary experimental results. In: 2nd International conference on multi-agent systems, AAAI Press, pp 95–102

Gerding E, Poutré H (2006) Bilateral bargaining with multiple opportunities: knowing your opponent's bargaining position. IEEE Trans Syst Man Cybern C 36:45–55
Gerding E, Bragt D, Poutré H (2003) Multi-issue negotiation processes by evolutionary simulation, validation and social extensions. Comput Econ 22:39–63
Giunchiglia F, Serafini L (1994) Multilanguage hierarchical logics, or: how we can do without modal logics. Artif Intell 65:29–70
Gordon T, Karacapilidis N (1997) The Zeno argumentation framework. In: 6th international conference on AI and law, ACM Press, New York, pp 10–18
Greenhalgh L (2001) Managing strategic relationships. Free Press, New York
He M, Jennings N, Leung H-F (2003) On agent-mediated electronic commerce. IEEE Trans Knowl Data Eng 15:985–1003
Heiskanen H, Ehtamo P, Hämäläinen R (2001) Constraint proposal method for computing Pareto solutions in multi-party negotiations. Eur J Oper Res 133:44–61
Holmes M (1992) Phase structures in negotiation. In: Putnam L, Roloff M (eds) Communication and negotiation. Sage, London pp 83–105
Jennings N, Sycara K, Wooldridge M (1998) A roadmap of agent research and development. Auton Agent Multi Agent Syst 1:275–306
Jennings N, Faratin P, Lomuscio A, Parsons S, Wooldridge M, Sierra C (2001) Automated negotiation: prospects, methods and challenges. Group Decis Negot 10:199–215
Jung H, Tambe M, Kulkarni S (2001) Argumentation as distributed constraint satisfaction. In: 5th International Conference on Autonomous Agents. ACM Press, New York, pp 324–331
Karunatillake N, Jennings N (2005) Is it worth arguing? In: 1st international workshop on argumentation in multi-agent systems. Springer, Heidelberg, pp 234–250 (LNAI 3366)
Kelley H (1966) A classroom study of the Dilemmas in interpersonal negotiation. In: Archibald K (ed) Strategic interaction and conflict: original papers and discussion. University of California Press, Berkeley, pp 49–73
Keough C (1992) Bargaining arguments and argumentative bargainers. In: Putnam L, Roloff M (eds) Communication and negotiation. Sage, London, pp 109–127
Kersten G, Lai H (2007) Negotiation support and e-negotiation systems: an overview. Group Decis Negot 16:553–586
Klein M (1991) Supporting conflict resolution in cooperative design systems. IEEE Trans Syst Man Cybern 21:1379–1390
Klein M, Faratin P, Sayama H, Bar-Yam Y (2003) Negotiating complex contracts. Group Decis Negot 12: 111–125
Kramer R (1991) The more the merrier? Social psychological aspects of multiparty negotiations in organizations. In: Bazerman M, Lewicki R, Sheppard B (eds) Research on negotiation in organizations, vol 3. JAI Press, London pp 307–332
Kraus S (2001) Strategic negotiation in multi-agent environments. MIT Press, Cambridge
Kraus S, Wilkenfeld J, Zlotkin G (1995) Multiagent negotiation under time constraints. Artif Intell 75:297–345
Kraus S, Sycara S, Evenchik A (1998) Reaching agreements through argumentation: a logical model and implementation. Artif Intell 104:1–69
Krause P, Ambler S, Elvang Gransson M, Fox J (1995) A logic of argumentation for reasoning under uncertainty. Comput Intell 11:113–131
Kressel K, Pruitt D (1989) Mediation research. Jossey-Bass, San Francisco
Lander S (1994) Distributed search and conflict management among reusable heterogeneous agents. Dissertation, University of Massachusetts at Amherst, USA
Lander S, Lesser V (1997) Sharing metainformation to guide cooperative search among heterogeneous reusable agents. IEEE Trans Knowl Data Eng 9:193–207
Lewicki R, Saunders D, Minton J (1999) Negotiation: readings, exercises, and cases. McGraw Hill, New York
Lewicki R, Barry B, Saunders D, Minton J (2003) Negotiation. McGraw Hill, New York
Li C, Giampapa J, Sycara K (2004) Bilateral negotiation decisions with uncertain dynamic outside options. In: 1st international workshop on electronic contracting, IEEE Computer Society Press, pp 54–61
Li C, Giampapa J, Sycara K (2006) Bilateral negotiation decisions with uncertain dynamic outside options. IEEE Trans Syst Man Cybern C 36:31–44
Loewenstein J, Thompson L, Gentner D (1999) Analogical encoding facilities knowledge transfer in negotiation. Psychon Bull Rev 6:586–597
Lomuscio A, Wooldridge M, Jennings N (2003) A classification scheme for negotiation in electronic commerce. Group Decis Negot 12:31–56
Lopes F, Mamede N, Novais AQ, Coelho H (2002) A negotiation model for autonomous computational agents: formal description and empirical evaluation. J Intell Fuzzy Syst 12:195–212

Lopes F, Mamede N, Novais AQ, Coelho H (2004) Negotiation strategies for autonomous computational agents. In: 16th European conference on artificial intelligence (ECAI-04). IOS Press, pp 38–42
Lopes F, Mamede N, Novais AQ, Coelho H (2005) Negotiation among autonomous agents: experimental evaluation of integrative strategies. In: 12th Portuguese conference on artificial intelligence. IEEE Computer Society Press, pp 280–288
Lovenheim P (1989) Mediate, don't litigate: how to resolve disputes quickly, privately, and inexpensively without going to court. McGraw Hill, New York
Maes P, Guttman R, Moukas A (1999) Agents that buy and sell. Commun ACM 42:81–91
Malsch T, Weiss G (2000) Conflicts in social theory and multiAgent systems: on importing sociological insights into distributed AI. In: Tessier C, Chaudron L, Müller H (eds) Conflicting agents: conflict management in multi-agent systems. Kluwer, Boston, pp 111–149
Matos N, Sierra C, Jennings N (1998) Determining successful negotiation strategies. In: 3rd international conference on multi-agent systems. IEEE Computer Society Press, pp 182–189
McGillicuddy N, Welton G, Pruitt D (1987) Third party intervention: a field experiment comparing three different models. J Person Soc Psychol 53:104–112
Moore C (1996) The mediation process: practical strategies for resolving conflict. Jossey-Bass, San Francisco
Morley I, Stephenson G (1977) The social psychology of bargaining. Allen and Unwin, London
Müller H (1996a) Negotiation principles. In: O'Hare G, Jennings N (eds) Foundations of distributed artificial intelligence. Wiley, Chichester, pp 211–299
Müller J (1996b) The design of intelligent agents. Springer, Heidelberg
Nadler J, Thompson L, Van Boven L (2003) Learning negotiation skills: four models of knowledge creation and transfer. Manag Sci 49:529–540
Nguyen T, Jennings N (2003) A heuristic model for concurrent bi-lateral negotiations in incomplete information settings. In: 18th international joint conference on artificial intelligence. Morgan Kaufmann Publishers, San Mateo, pp 1467–1469
Nguyen T, Jennings N (2004) Coordinating multiple concurrent negotiations. In: 3rd international conference on autonomous agents and multi agent systems. ACM Press, New York, pp 1062–1069
Nguyen T, Jennings N (2005) Managing commitments in multiple concurrent negotiations. Electron Commer Res Appl 4:362–376
Olekalns M, Smith P, Walsh T (1996) The process of negotiating: strategy and timing as predictors of outcomes. Organ Behav Hum Decis Process 68:68–77
Oliver J (1997) A machine-learning approach to automated negotiation and prospects for electronic commerce. J Manag Inf Syst 13:83–112
Osborne M, Rubinstein A (1990) Bargaining and markets. Academic Press, San Diego
Osborne M (2004) Introduction to game theory. Oxford University Press, Oxford
Parsons S, Jennings N (1996) Negotiation through argumentation—a preliminary report. In: 2nd International Conference on Multi-Agent Systems. AAAI Press, pp 267–274
Parsons S, McBurney P (2003) Argumentation-based dialogues for agent co-ordination. Group Decis Negot 12:415–439
Parsons S, Sierra C, Jennings N (1998) Agents that reason and negotiate by arguing. J Log Comput 8:261–292
Prasad M, Lesser V, Lander S (1998) Learning organizational roles for negotiated search in a multiagent system. Int J Hum Comp Stud 48:51–67
Prakken H, Vreeswijk G (2002) Logics for defeasible argumentation. In: Gabbay D, Guenthner F (eds) Handbook of philosophical logic. Kluwer, Boston, pp 218–319
Pruitt D (1981) Negotiation behavior. Academic Press, New York
Pruitt D (1998) Social conflict. In: Gilbert D, Fiske S, Lindzei G (eds) The handbook of social psychology, vol 2. McGraw-Hill, New York, pp 470–503
Pruitt D, Rubin J (1986) Social conflict: escalation, stalemate and settlement. McGraw Hill, New York
Pruitt D, Carnevale P (1993) Negotiation in social conflict. Open University Press, Philadelphia
Pruitt D, Kim S (2004) Social conflict: escalation, stalemate, and settlement. McGraw Hill, New York
Putnam L (1990) Reframing integrative and distributive bargaining. In: Shepard B, Bazerman M, Lewicki R (eds) Research on negotiation in organizations, vol 2. JAI Press, London, pp 3–30
Putnam L, Wilson S, Turner D (1990) The evolution of policy arguments in teacher's negotiations. Argumentation 4:129–152
Rahwan I, Ramchurn S, Jennings N, McBurney P, Parsons S, Sonenberg L (2004) Argumentation-based negotiation. Knowl Eng Rev 18:343–375
Raiffa H (1982) The art and science of negotiation. Harvard University Press, Cambridge
Raiffa H (1985) Post-settlements settlements. Negot J 1:9–12
Raiffa H (2002) Negotiation analysis. Harvard University Press, Cambridge

Reed C, Norman T (2004) Argumentation machines—new frontiers in argument and computation. Kluwer, Boston
Roloff M, Jordan J (1992) Achieving negotiation goals: the fruits and foibles of planning ahead. In: Putnam L, Roloff M (eds) Communication and negotiation. Sage, London, pp 21–45
Rosenschein J, Zlotkin G (1994) Rules of encounter: designing conventions for automated negotiation among omputers. MIT Press, Cambridge
Ross B, Kilbane M (1997) Effects of principle explanation and superficial similarity on analogical mapping in problem solving. J Exp Psychol Learn Mem Cogn 23:427–440
Sandholm T (1999) Distributed rational decision making. In: Weiss G (ed) MultiAgent systems—a modern approach to distributed artificial intelligence. MIT Press, Cambridge, pp 201–259
Sandholm T, Lesser V (2001) Leveled-commitment contracts and strategic breach. Games Econ Behav 35:212–270
Sandholm T, Lesser V (2002) Leveled-commitment contracting: a backtracking instrument for multi-agent systems. AI Mag 23:89–100
Sabater J, Sierra C, Parsons S, Jennings N (2002) Engineering executable agents using multi-context systems. J Log Comput 12:413–442
Sathi A, Fox M (1989) Constraint-directed negotiation of resource reallocations. In: Gasser L, Huhns M (eds) Distributed artificial intelligence II. Morgan Kaufmann Publishers, San Mateo, pp 163–193
Savage G, Blair J, Sorenson R (1989) Consider both relationships and substance when negotiating strategically. Acad Manag Exec 3:37–48
Sen S, Weiss G (1999) Learning in multiagent systems. In: Weiss G (ed) MultiAgent systems—a modern approach to distributed artificial intelligence. MIT Press, Cambridge pp 259–298
Sierra C, Jennings N, Noriega P, Parsons S (1998) A framework for argumentation-based negotiation. In: Intelligent agents IV—international Workshop on agent Theories, architectures and languages. Springer, Heidelberg, pp 177–192 (LNAI 1365)
Shintani T, Ito T, Sycara K (2000) Multiple negotiations among agents for a distributed meeting scheduler. In: 4th International Conference on MultiAgent Systems. IEEE Computer Society Press, pp 435–436
Simon H (1981) The sciences of the artificial. MIT Press, Cambridge
Smith R (1980) The contract net protocol: high-level communication and control in a distributed problem solver. IEEE Trans Comput 29:1104–1113
Snyder G, Diesing P (1977) Conflict among nations. Princeton University Press, Princeton
Sycara K (1987) Resolving adversarial conflicts: an approach integrating case-based and analytic methods. Dissertation, School of Information and Computer Science, Georgia
Sycara K (1990) Persuasive argumentation in negotiation. Theory Decis 28:203–242
Sycara K (1991) Problem restructuring in negotiation. Manag Sci 37:1248–1268
Sycara K (1993) Machine learning for intelligent support of conflict resolution. Decis Support Syst 10: 121–136
Sycara K (1998) Multi-agent systems. AI Mag 19:79–92
Thompson L (1991) Information exchange in negotiation. J Exp Soc Psychol 27:161–179
Thompson L (2005) The mind and heart of the negotiator. Prentice-Hall, Englewood Cliffs
Thompson L, DeHarpport T (1994) Social judgment, feedback, and interpersonal learning in negotiation. Org Behav Hum Decis Process 58:327–345
Thompson L, Gentner D, Loewenstein J (2000) Avoiding missed opportunities in managerial life: analogical training more powerful than individual case training. Org Behav Hum Decis Process 82:60–75
Toulmin S, Rieke R, Janik A (1979) An introduction to reasoning. MacMillan, London
Tutzauer F (1992) The communication of offers in dyadic bargaining. In: Putnam L, Roloff M (eds) Communication and negotiation. Sage, London, pp 67–82
Vande Vliert E (1997) Complex interpersonal conflict behavior: theoretical frontiers. Psychology Press, East Sussex
Von Martial F (1992) Coordinating plans of autonomous agents. Springer, Heidelberg
Wagner T, Shapiro J, Xuan P, Lesser V (1999) Multi-level conflict in multi-agent systems. In: AAAI-99 Workshop on Negotiation in Multi-agent Systems. AAAI Press, pp 50–55
Walton D, Krabbe C (1995) Commitment in dialogue: basic concepts of interpersonal reasoning. State University of New York Press, New York
Weingart L, Thompson L, Bazerman M, Carroll J (1990) Tactical behaviors and negotiation outcomes. Int J Confl Manag 1:7–31
Weingart L, Bennett R, Brett J (1993) The impact of consideration of issues and motivational orientation on group negotiation process and outcome. J Appl Psychol 78:504–517
Weiss G (1999) MultiAgent systems—a modern approach to distributed artificial intelligence. MIT Press, Cambridge

Wooldridge M (2002) An introduction to multiAgent systems. Wiley, Chichester
Wurman P, Wellman M, Walsh W (2001) A parametrization of the auction design space. Games Econ Behav 35:304–338
Yokoo M, Hirayama H (1998) Distributed constraint satisfaction algorithm for complex local problems. In: 3rd International Conference on Multi-Agent Systems. IEEE Computer Society Press, pp 372–379
Zadeh I (1971) Similarity relations and fuzzy orderings. Inf Sci 3:177–200
Zartman I, Berman M (1982) The practical negotiator. Yale University Press, New Haven
Zeng D, Sycara K (1998) Bayesian learning in negotiation. Int J Hum Comput Stud 48:125–141
Zhang X, Lesser V, Abdallah S (2005) Efficient management of multi-linked negotiation based on a formalized model. Auton Agent Multi Agent Syst 10:165–205
Zubek J, Pruitt D, Peirce R, McGillicuddy N, Syna H (1992) Disputant and mediator behaviors affecting short-term success in mediation. J Confl Resolut 36:546–572