Journal of Intelligent & Fuzzy Systems 12 (2002) 195–212
IOS Press

A negotiation model for autonomous computational agents: Formal description and empirical evaluation

Fernando Lopes[a,*], Nuno Mamede[b], A.Q. Novais[a] and Helder Coelho[c]
[a]*Departamento Modelação e Simulação de Processos, INETI, Est. Paço Lumiar, 1649-038 Lisboa, Portugal*
[b]*Departamento Engenharia Informática, IST, Av. Rovisco Pais, 1049-001 Lisboa, Portugal*
[c]*Faculdade Ciências, Departamento Informática, UL, Campo Grande, 1749-016 Lisboa, Portugal*

Abstract. Autonomous agents are being used in an increasing number of applications. The agents operate in complex environments and, over time, conflicts inevitably occur among them. Negotiation is the predominant process for resolving conflicts. This paper presents a generic negotiation model for autonomous agents that handles multi-party, multi-issue and single or repeated rounds. The model is based on computationally tractable assumptions and accounts for a tight integration of the individual capability of planning and the social capability of negotiation. This paper also describes an experiment conducted to evaluate the model in different types of situations. The experimental results confirmed a number of well-documented conclusions about human negotiation.

1. Introduction

Autonomous software agents are being used in an increasing number of applications [20]. These agents have the ability to decide for themselves which goals to adopt, which actions to perform in order to achieve these goals, and when to perform these actions. Most applications involve or require multiple agents operating in complex environments and, over time, conflicts inevitably occur among them. Conflict resolution is crucial for achieving multi-agent coordination. The predominant process for resolving conflicts is negotiation – the process by which two or more agents attempt to influence other agents in an effort to achieve their needs, while at the same time taking the needs of the others into account [23].

Artificial intelligence (AI) researchers have recently started to investigate the design of autonomous negotiating agents (e.g. [19,42]). Some researchers developed or adopted a model of individual behavior and used the model as a starting point for the development of a negotiation model. However, most researchers have focused solely on developing negotiation models. They have addressed only part of the overall task of building autonomous negotiating agents. In particular, they have paid little attention to the problem of integrating existing or new models of individual behavior with their negotiation models. This fundamental problem is still an open problem.

This paper presents a generic negotiation model for autonomous agents that handles multi-party, multi-issue, and single or repeated rounds. The main components of the model are: (i) a prenegotiation model, (ii) a multilateral and a bilateral negotiation protocols, (iii) an individual model of the negotiation process, (iv) a set of negotiation strategies, and (v) a set of negotiation tactics. The model is based on computationally tractable assumptions, accounts for a tight integration of the individual capability of planning and the social capability of negotiation, and formalizes a set of human negotiation procedures.

The model is currently being evaluated. This paper presents a detailed description of an experiment

*Corresponding author. E-mail: fernando.lopes@ineti.pt.

conducted to: (i) assess the feasibility of building autonomous negotiating agents equipped with a simplified version of the model, (ii) investigate the integration of planning and negotiation, and (iii) evaluate the effect of different strategies both on the convergence of the negotiation process and on the outcome of negotiation. The experimental results confirmed a number of well-documented conclusions about human negotiation.

This paper builds on our previous work in the area of negotiation. In particular, it extends the prenegotiation model and the individual model of the negotiation process presented in [24–26]. It also extends the set of negotiation strategies and tactics presented in [27,28]. The work described here is also complementary to the work described in these papers, because it concentrates both on the negotiation model and the empirical evaluation of the model rather than on the theoretical model alone. Finally, this paper fixes a few technical problems associated with the components of the model described in these papers.

The remainder of the paper is structured as follows. Section 2 presents the main approaches followed by AI researchers for developing autonomous negotiating agents. This section places our work in the context of previous work. Section 3 presents a generic model of individual behavior for autonomous agents and formalizes the concept of conflict of interest. The work described in this section is the starting point for our work. Section 4 presents a generic model of negotiation for autonomous agents. Section 5 describes the experimental evaluation of the negotiation model. Section 6 compares the negotiation model with other developed models. Finally, Section 7 concludes and outlines a number of issues which require further investigation.

2. The design of autonomous negotiating agents

The design of autonomous agents with negotiation competence has been investigated by AI researchers from both a theoretical and a practical perspective.

Researchers following the theoretical perspective attempt mainly to develop formal models. Some researchers define the modalities of the mental state of the agents (e.g., beliefs, desires and intentions), develop a *logical* model of individual behavior, and then use the model as a basis for the development of a formal model of negotiation or argumentation (e.g. [19]). However, most researchers are neutral with respect to the modalities of the mental state and just develop formal models of negotiation. These models are often based on game-theoretical techniques (e.g. [18,39]).

Generally speaking, most theoretical models are rich but restrictive. They make a number of assumptions that severely limit their applicability to solve real problems. To a large extent, they are not concerned with computational issues. As a result, they require substantial computational effort.

Researchers following the practical perspective attempt mainly to develop *computational* models, i.e., models specifying the key data structures of the agents and the processes operating on these structures. Again, some researchers start with a particular model of individual behavior (e.g., a belief-desire-intention model), develop a negotiation model or adopt an existing one, and then integrate both models into a unified model that accounts for both individual and social behavior (e.g. [30]). However, most researchers prefer to be neutral about the model of individual behavior and just develop models of negotiation (e.g. [7,42]).

Broadly speaking, most computational models are rich but lack a rigorous theoretical grounding. As a result, there is no precise understanding of how the computer systems resulting from these models work in the way they do.

This work seeks to develop autonomous negotiating agents for operating in complex application domains (e.g., a supply chain). As noted, both the theoretical and the practical perspectives have specific strengths and weaknesses. However, despite the weaknesses of the practical perspective, some researchers believe that it is necessary to develop computational models in order to implement and successfully use autonomous agents in real-world applications [38]. Accordingly, this paper presents a computational model of negotiation. Also, as noted, most researchers following the practical perspective have paid little attention to the problem of integrating models of individual behavior and negotiation models. However, it is one of the commonest and costliest lessons of computer science that independently developed components resist subsequent integration in a smoothly functioning whole. Components need to be designed for integration right from the start [14]. Accordingly, this paper presents a model that accounts for a tight integration of the individual capability of planning and the social capability of negotiation.

As a last point, most researchers following the practical perspective have paid little attention to a number of issues. We highlight the following ones:

1. What is a conflict? How do agents acknowledge the role of conflict as a driving force of negotiation?
2. How to plan and prepare for negotiation? Which are the activities that agents must attend to before actually starting to negotiate?
3. What is a negotiation problem? How do agents represent negotiation problems?
4. How do agents determine the set of negotiation issues?
5. What are negotiation strategies? How they are formalized? Are they based on human negotiation procedures?
6. How can agents change the representation of negotiation problems? How can they dynamically add and remove negotiation issues?

This paper addresses these issues in a domain-independent way.

3. Autonomous agents and conflict of interests

The first part of this section presents a generic model of individual behavior for autonomous agents. This statement requires some qualification, however. Even a superficial reading of the literature demonstrates the existence of a wide range of agents – different researchers have different ideas about what agents are. Therefore, the model is not a canonical model of autonomous agents. Also, the model is not a complete model of autonomous agents. The aim is to present a computational model that captures some of the most important features of a wide range of agents.

The second and last part of this section defines formally the concept of conflict of interests, presents axioms for conflict detection, and describes a procedure for conflict validation.

The work described here forms a basis for the development of autonomous negotiating agents. It is the starting point for our work.

3.1. Autonomous agents

Let $Ag = \{ag_1, \ldots, ag_n\}$ be a set of autonomous agents. A description of the key features of every agent $ag_i \in Ag$ follows.

Beliefs, Goals and Plan Templates. The agent ag_i has a set $B_i = \{b_{i1}, b_{i2}, \ldots\}$ of beliefs, a set $G_i = \{g_{i1}, g_{i2}, \ldots\}$ of goals, and a library $PL_i = \{pt_{i1}, pt_{i2}, \ldots\}$ of plan templates.

Beliefs represent information about the world and the agent himself. Goals represent world states to be achieved. Plan templates are simple procedures for achieving goals. Every plan template $pt_{ij} \in PL_i$ is a 6-tuple:

$$pt_{ij} = < header_{ij},\ type_{ij},\ preconds_{ij},\ body_{ij}, constrs_{ij}, effects_{ij} >$$

The header is a 2-tuple: $header_{ij} =< name_{ij}, vars_{ij} >$, where $name_{ij}$ is the name of pt_{ij} and $vars_{ij}$ is a set of variables (arguments of pt_{ij}). In most cases, the header is simply the description of a goal $g_{ij} \in G_i$ for which pt_{ij} is a recipe. The $type_{ij}$ is the type of pt_{ij} (composite or primitive). $Preconds_{ij}$ is a list of conditions that must hold before pt_{ij} can be applied. The $body_{ij}$ is either a list of subgoals whose achievement constitutes the achievement of a goal g_{ij} or a list of primitive actions (i.e., actions directly executable by ag_i) whose performance constitutes the achievement of g_{ij}. $Constrs_{ij}$ is a list of constraints (e.g., to impose a temporal order on the members of the body). $Effects_{ij}$ is a list of statements that hold after pt_{ij} has been successfully executed.

The library PL_i has composite and primitive plan templates. A $composite$ plan template is a recipe specifying the decomposition of a goal into a set of subgoals. A $primitive$ plan template is a recipe specifying a primitive action or a sequence of primitive actions that can achieve a goal.

Plan Generation. The agent ag_i is able to generate complex plans from the simpler plan templates stored in the library.

A $plan$ p_{ik} for achieving a goal $g_{ik} \in G_i$ is a 3-tuple:

$$p_{ik} =< PT_{ik}, \leqslant_h, \leqslant_t >$$

where $PT_{ik} \subseteq PL_i$ is a list of instantiated plan templates (i.e., plan templates where some or all of the *arguments* have been instantiated), $\leqslant_h$ is a binary relation establishing a hierarchy on PT_{ik} ($pt_{ik1} \leqslant_h pt_{ik2}$, for $pt_{ik1} \in PT_{ik}$ and $pt_{ik2} \in PT_{ik}$, means that pt_{ik2} is an immediate successor of pt_{ik1}, i.e., a successor for which no intermediate plan templates are permitted), and $\leqslant_t$ is another binary relation establishing a temporal order on PT_{ik} ($pt_{ik1} \leqslant_t pt_{ik2}$ means that pt_{ik1} must be applied before pt_{ik2}).

The plan p_{ik} is represented by a hierarchical and temporally constrained And-tree denoted by $Pstruct_{ik}$. The nodes of the tree are instantiated plan templates. Arcs form a hierarchy between pairs of nodes. Also, arcs represent ordering constraints.

The generation of p_{ik} is performed through an iterative procedure involving four main tasks: (i) plan retrieval, (ii) plan selection (iii) plan addition, and (iv) plan interpretation. These tasks are common to a wide range of hierarchical planning algorithms (see, for example [6,8,29]). A description of each task follows.

Plan retrieval consists of searching the plan library PL_i for any plan template whose header unify with the description of g_{ik} and retrieving all the plan templates $AP_{ik} = \{pt_{ik1}, pt_{ik2}, \ldots, pt_{ikp-1}, pt_{ikp}, pt_{ikp+1}, \ldots, pt_{ikz}\}$ whose preconditions hold in the current state (i.e., the preconditions are a logical consequence of the belief set B_i of ag_i). The plan templates in AP_{ik} are called *applicable* plan templates.

Plan selection consists of selecting the preferred plan template $pt_{ikp} \in AP_{ik}$. The plan templates in AP_{ik} are first evaluated by computing their score and then the plan template with the highest score is selected (see, for example [15,30]).

Plan addition consists of adding the selected plan template pt_{ikp} to p_{ik} and recording the remaining plan templates $RAP_{ik} = \{pt_{ik1}, pt_{ik2}, \ldots, pt_{ikp-1}, pt_{ikp+1}, \ldots, pt_{ikz}\}$ in p_{ik}. The plan templates in RAP_{ik} are called alternative plan templates and have a key role in the definition of a structure for a negotiation problem (see Subsection 4.1). They are explicitly recorded in p_{ik} and placed alongside pt_{ikp}.

Plan interpretation consists of selecting a composite plan template from p_{ik}, say pt_{ikp}, establishing a temporal order for the elements of the $body_{ikp} = [g_{ikp+1}, g_{ikp+2}, \ldots]$ of pt_{ikp}, and selecting the first ordered element g_{ikp+1}. The temporal order is defined by the list of constraints $constrs_{ikp}$. The elements of $body_{ikp}$ are interpreted as subgoals of the goal g_{ik}.

Adopted Plans. At any instant, the agent ag_i has a number of plans for execution, either immediately or in the near future. These plans are the plans *adopted* by ag_i and are stored in the *intention structure* IS_i. Formally, IS_i is defined as follows:

$$IS_i = [p_{i1}, p_{i2}, \ldots, p_{ik}, \ldots]$$

For each plan $p_{im} \in IS_i$, the header of every plan template pt_{imj} in p_{im} is referred as *intention* int_{imj}. Intentions are therefore goals not yet achieved and considered achievable – goals restricted to the existence of plans for achieving them.

It is worth noting that the term "adopted plan" entails a commitment to act in order to satisfy, or attempt to satisfy, the intentions that constitute a plan. The nature of this commitment is quite complex (see, for example [2,10,33]). However, this commitment means at least that the plans adopted by an agent should be reasonably stable, i.e., they should be subject to reconsideration only at appropriate (crucial) moments. This raises the important and hard question of when to reconsider the adopted plans. To simplify matters in this respect, we consider that an agent commits to the plans he adopts and undertakes to change them only when they conflict with the plans of other agents. In particular, the agents negotiate mutually acceptable agreements that often lead to plan reconsideration.

Social Description. The agent ag_i often has information about the other agents in Ag. This information can be acquired either through perception or communication and is stored in the *social description* SD_i. Formally, SD_i is defined as follows:

$$SD_i = \{SD_i(ag_1), SD_i(ag_2), \ldots, SD_i(ag_n)\}$$

where each structure $SD_i(ag_j) \in SD_i$ holds information about a particular agent $ag_j \in Ag$. More specifically, each structure is a 3-tuple:

$$SD_i(ag_j) =< B_i(ag_j), G_i(ag_j), I_i(ag_j) >$$

where $B_i(ag_j), G_i(ag_j)$, and $I_i(ag_j)$ are the sets of beliefs, goals and intentions that ag_i believes ag_j has, respectively.

The information in SD_i may be both incomplete and incorrect. Incompleteness means that some information is missing (e.g., ag_i believes that ag_j has formulated a plan p_{jk} but has only information about a few intentions included in p_{jk}). Incorrectness means that some information is outdated.

3.2. Conflict of interests

Let $ag_i \in Ag$ be an agent with a plan p_{ik} including intention int_{ikp}. Let $A = \{ag_1, \ldots, ag_n\}$, $A = Ag - \{ag_i\}$, be a set of agents that interact with ag_i. Let IS_i be the intention structure of ag_i and $SD_i = [SD_i(ag_1), \ldots, SD_i(ag_n)]$ be his social description.

Let $PP = \{p_{i1}(ag_1), \ldots, p_{in}(ag_n)\}$ be a set of *possible plans* of the agents in A, i.e., plans that ag_i believes these agents have generated. Let $PI = \{int_{i11}(ag_1), \ldots, int_{inn}(ag_n)\}$ be a set of *possible intentions* of the agents in A, i.e., intentions that ag_i believes these agents have formulated as part of plans $\{p_{i1}(ag_1), \ldots, p_{in}(ag_n)\}$, respectively.

Let the intentions in $PI \cup \{int_{ikp}\}$ represent commitments to achieve exclusive world states. In this situation, the intentions are called *incompatible* and represented by $Incomp(int_{ikp}, int_{i11}(ag_1), \ldots,$

$int_{inn}(ag_n))$, emphasizing the fact that they cannot be executed together. The plans in $PP \cup \{p_{ik}\}$ are also called *incompatible* and represented by $Incomp(p_{ik}, p_{i1}(ag_1), \ldots, p_{in}(ag_n))$.

Potential Conflict of Interests. A *potential* conflict of interests from the perspective of ag_i and with respect to plan p_{ik} (intention int_{ikp}) is defined formally as follows (see, for example [5,9,41]):

$$PotConf_{ik} =$$
$$\exists int_{ikp} \in IS_i \wedge \exists int_{i11}(ag_1) \in SD_i(ag_1)$$
$$\wedge \ldots \wedge int_{inn}(ag_n) \in SD_i(ag_n) \wedge$$
$$Incomp(int_{ikp}, int_{i11}(ag_1), \ldots, int_{inn}(ag_n))$$

It is important to note that potential conflict is defined as being subjective, i.e., ag_i only needs to believe the agents in A intend to achieve specific world states, and does not need to know the real intentions of these agents.

Potential Conflict Detection. The agents in Ag check regularly their adopted plans in order to detect any potential conflict of interests. Conflict detection is done individually by each agent $ag_i \in Ag$. To this end, ag_i has a library of conflict detection axioms $CL_i = \{ax_{i1}, ax_{i2}, \ldots\}$. Every axiom $ax_{ik} \in CL_i$ has the following generic form:

$$int_{ikp} \& int_{i11}(ag_1) \& \ldots \& int_{inn}(ag_n) \&$$
$$conds \rightarrow false$$

where $int_{ikp}, int_{i11}(ag_1)$ and $int_{inn}(ag_n)$ have the meaning just specified, $conds$ is a list of conditions, *false* is a 0-ary predicate symbol, & is the conjunction operator, and $\rightarrow$ the implication operator. The axiom ax_{ik} states that the intentions $(int_{ikp}, int_{i11}(ag_1), \ldots, int_{inn}(ag_n))$ represent commitments to achieve exclusive world states and, therefore, cannot be satisfied together.

Potential Conflict Validation. Potential Conflict validation is a process by which the conflicting agents in Ag carry out a conversation towards the goal of confirming the possible intentions used in conflict detection.

Let ag_i be an agent that detects a potential conflict of interests $PotConf_{ik}$ ($Conf$, for short). There are many different conversations the agents in Ag may carry out to fulfill the goal of confirming the intentions in PI. A specific conversation taking place between ag_i on one side and every agent in A on the other side follows (see, for example [30]).

The conversation starts with ag_i announcing the detection of the potential conflict. This is done by sending an inform message containing the conflict identifier $Conf$. Every agent in A can either: (i) decide to discuss the nature of the conflict or (ii) do nothing. The former decision leads to an acknowledgement of the inform message sent by ag_i. The latter decision results in a timeout and ends the conversation. If all agents in A acknowledge the inform message, then ag_i requests them to inform whether the information used in the detection of $Conf$ is true. More specifically, ag_i sends to each agent $ag_j \in A$ a request message to inform about the truthfulness of a possible intention $int_{ijj}(ag_j), 1 \leqslant j \leqslant n, j \neq i$.

Upon receiving the request, the agents in A have the choice of either: (i) confirming or not the possible intentions, or (ii) doing nothing. In the first case, every agent ag_j sends to ag_i an inform message containing either $int_{ijj}(ag_j)$ or $\neg(int_{ijj}(ag_j))$, where $\neg$ is the negation operator. The confirmation of all the possible intentions in PI results in the validation of the conflict. In the second case, if at least one agent decides to do nothing, the conversation ends. The agent ag_i receives all the inform messages and based on their number and content decides either to validate or not validate the conflict. The former decision is followed by conflict declaration. This is done by ag_i sending a declare message containing $Conf$, the intention int_{ikp}, the set A of agents, and the set PI of (confirmed) intentions. The latter decision leads to ag_i sending a declare message containing $\neg(Conf)$. The conversation ends with the agents in A acknowledging the declare message.

This conversation exhibits two desirable aspects. First, it is intuitive and to a certain extent corresponds to the way humans validate information. Second, it is simple, requiring little communication overhead and consuming few computational resources. However, this conversation lacks both symmetric distribution and generality. In fact, ag_i plays a central role – he initiates the conversation, communicates with each one of the other agents, reasons about the feedback received from these agents, and decides about conflict validation. Also, the conversation is only appropriate for agents that are willing to reveal their intentions truthfully without compensation, if asked by other agents.

The validation of potential conflicts of interests leads to *true* conflicts of interests (hereafter, just referred as conflicts).

4. The negotiation model

Negotiation is the predominant process for resolving conflicts. Examination of the literature in the fields of social psychology (e.g. [3,34–36]), economy and game theory (e.g. [22,23,31,37]), and distributed artificial intelligence (e.g. [7,30,32,39]) motivated the development of a generic negotiation model that handles multi-party, multi-issue, and single or repeated rounds. The main components of the model are:

1. a prenegotiation model;
2. a bilateral and a multilateral negotiation protocols;
3. an individual model of the negotiation process;
4. a set of negotiation strategies;
5. a set of negotiation tactics.

This section presents a domain-independent and formal description of each component.

4.1. Preparing and planning for negotiation

Successful negotiators agree on one thing: the key to success in negotiation is preparation and planning. Persuasive presentation, skillful communication, and a host of other skills used during negotiation are important, but they cannot overcome the disadvantage created by a poor planning [23].

The prenegotiation model defines the main activities that each agent $ag_i \in Ag$ must attend to in order to prepare and plan for negotiation. A formal description of each activity follows.

Negotiation Problem Definition and Structure Generation. Conflicts raise negotiation problems. Formally, a *negotiation problem* from the perspective of ag_i is a 7-tuple:

$$NP_{ik} = < ag_i, B_i, g_{ik}, p_{ik}, int_{ikp}, A, I_A >$$

where $B_i, g_{ik}, p_{ik}, int_{ikp}$ and A have the meaning just specified, and I_A is a set of intentions of the agents in A incompatible with intention int_{ikp}.

The problem NP_{ik} has a $structure\ NPstruct_{ik}$ consisting of a hierarchical And-Or tree. Formally, $NPstruct_{ik}$ is a 4-tuple:

$$NPstruct_{ik} = < NPT_{ik}, \leqslant_h, \leqslant_t, \leqslant_a >$$

where $NPT_{ik} \subseteq PL_i$ is a list of instantiated plan templates, $\leqslant_h$ and $\leqslant_t$ have the meaning specified in Subsection 3.1, and $\leqslant_a$ is a binary relation establishing alternatives among the plan templates in $NPT_{ik}(pt_{ik1} \leqslant_a pt_{ik2}$, for $pt_{ik1} \in NPT_{ik}$ and $pt_{ik2} \in NPT_{ik}$, means that pt_{ik1} and pt_{ik2} are alternative ways for achieving the goal specified by the header of either plan templates). The nodes of the And-Or tree are plan templates. The header of the root node describes a $negotiation\ goal\ g_{ik}$.

The structure $NPstruct_{ik}$ of NP_{ik} is generated from plan p_{ik}. First, an initial structure is generated for NP_{ik}. This structure is simply a copy of p_{ik}'s structure (And tree). Next, the plan p_{ik} is expanded through an iterative procedure involving the following tasks: (i) plan interpretation, (ii) plan retrieval, (iii) plan selection, and (iv) plan addition. These tasks were described in Subsection 3.1 and, for this reason, are only summarized below.

Plan interpretation consists of selecting an *alternative* plan template pt_{ikc} from the structure of NP_{ik}, establishing a temporal order for the elements of the $body_{ikc} = [g_{ikc+1}, g_{ikc+2}, \ldots]$ of pt_{ikc}, and selecting the first ordered element g_{ikc+1}. *Plan retrieval* consists of searching the plan library PL_i and finding all the plan templates $NAP_{ik} = \{pt_{ik1}, \ldots, pt_{ikl-1}, pt_{ikl}, pt_{ikl+1}, \ldots, pt_{ikz}\}$ whose name and arguments match the description of g_{ikc+1}. *Plan selection* consists of arbitrarily selecting a plan template $pt_{ikl} \in NAP_{ik}$. *Plan addition* consists of adding the selected plan template pt_{ikl} to the plan p_{ik} and recording the remaining plans $RNAP_{ik} = \{pt_{ik1}, \ldots, pt_{ikl-1}, pt_{ikl+1}, \ldots, pt_{ikz}\}$ in p_{ik}.

The complete expansion of the plan p_{ik} leads to $NPstruct_{ik}$. It is worth pointing out that $NPstruct_{ik}$ defines all the solutions of NP_{ik} currently known by ag_i. A *solution* is a plan that can achieve the negotiation goal g_{ik}.

Issue Identification and Prioritization. The negotiation issues of ag_i are obtained from the leaves of $NPstruct_{ik}$. Let $L_{ik} = [pt_{ika}, \ldots, pt_{ikz}, \ldots, pt_{ikz+n}]$ be the collection of primitive plan templates constituting the leaves of $NPstruct_{ik}$. The header ($name_{ikj}$ and $vars_{ikj}$) of every plan template $pt_{ikj} \in L_{ik}$ is called a $fact$ and denoted by f_{ikj}. Formally, a $fact\ f_{ikj}$ is a 3-tuple:

$$f_{ikj} = < is_{ikj}, v[is_{ikj}], r_{ikj} >$$

where is_{ikj} is a *negotiation issue* (corresponding to $name_{ikj}$), $v[is_{ikj}]$ is a *value* of is_{ikj} (corresponding to an argument of the list $vars_{ikj}$), and r_{ikj} is a list of arguments (corresponding to the remaining arguments of $vars_{ikj}$). Typically, r_{ikj} is an empty list.

Let $F_{ik} = \{f_{ika}, \ldots, f_{ikz}\}$ be the set of facts of $Npstruct_{ik}$ (F_{ik} has no duplicate facts). The *negotiating agenda* of ag_i is the set of issues $I_{ik} = \{is_{ika}, \ldots, is_{ikz}\}$ associated with the facts in F_{ik} (for clarity, we consider that every fact in F_{ik} is associated with a different issue).

The issues in I_{ik} can be either quantitative or qualitative. Quantitative issues are defined over continuous intervals. The interval of legal values for each quantitative issue $is_{ikq} \in I_{ik}$ is represented by $D_{ikq} = [min_{ikq}, max_{ikq}]$. Qualitative issues are defined over finite sets of values. The set of possible values for each qualitative issue $is_{ikx} \in I_{ik}$ is represented by $D_{ikx} = \{q_{ikx1}, q_{ikx2}, \ldots\}$.

The issues in I_{ik} are prioritized and ordered in a strictly descending order of preference. The *priority* of each issue $is_{ikj} \in I_{ik}$ is a number that represents its order of preference. The *weight* of is_{ikj} is a number that represents its relative importance. The sets of priorities and weights of the issues in I_{ik} are represented by $PR_{ik} = \{pr_{ika}, \ldots, pr_{ikz}\}$ and $W_{ik} = \{w_{ika}, \ldots, w_{ikz}\}$, respectively. The weights are normalized.

Limits and Aspirations Formulation. A *limit* or *reservation value* is a bargainer's ultimate fallback position, the level of benefit beyond which he is unwilling to concede. An *aspiration* is a level of benefit sought at any particular time, i.e., a value to the bargainer of the goal towards which he is striving. Limit tends to remain constant over time, whereas aspiration declines towards limit [34].

The agent ag_i formulates limits and aspirations for each issue $is_{ikj} \in I_{ik}$ at stake in negotiation. Let $T = \{t_1, t_2, \ldots\}$ be a linearly ordered set of instants representing the time. The *limit* for is_{ikj} is denoted by lim_{ikj} and the initial *aspiration* by asp^{t1}_{ikj}, with $lim_{ikj}, asp^{t1}_{ikj} \in D_{ikj}$.

Negotiation Constraints Definition. Negotiation constraints bound the possible values for the issues in I_{ik}. *Hard constraints* are linear boundary constraints that specify threshold values for the issues. They cannot be relaxed. *Soft constraints* are linear boundary constraints that specify minimum acceptable values for the issues. They can be relaxed, if necessary. They also can have different degrees of flexibility.

The agent ag_i defines constraints for each issue is_{ikj} in I_{ik}. Without loss of generality, consider that ag_i wants to maximize is_{ikj}. The hard constraint hc_{ikj} for is_{ikj} has the generic form:

$$hc_{ikj} = (is_{ikj} \geqslant lim_{ikj}, flex = 0)$$

where $flex = 0$ represents null flexibility (inflexibility). The soft constraint sc_{ikj} for is_{ikj} has the following similar form:

$$sc_{ikj} = (is_{ikj} \geqslant asp^{t1}_{ikj}, flex = n)$$

where asp^{t1}_{ikj} has the meaning just specified and $flex = n,\ n \in N$, represents the degree of flexibility of sc_{ikj}.

Negotiation Strategy Selection. The agent ag_i has a library $SL_i = \{str_{i1}, str_{i2}, \ldots\}$ of negotiation strategies and a library $TL_i = \{tact_{i1}, tact_{i2}, \ldots\}$ of negotiation tactics. *Negotiation strategies* are functions that define the tactics to be used at the beginning and during the course of negotiation (see Subsection 4.4). *Negotiation tactics* are functions that define the actions or moves to be made at each point of the negotiation process (see Subsection 4.5).

Strategy selection is an important task and must be carefully planned (see, for example [23,36,37]). The strategy most suitable for a particular negotiation situation often depends on the situation itself and cannot be specified in advance. As a result, strategy selection is a difficult task. In this paper, we assume that ag_i selects a strategy $str_{ik} \in SL_i$ that he considers appropriate according to his experience.

4.2. Negotiation protocols

The application of autonomous agents in areas such as electronic commerce has given increased importance to bilateral negotiation. Accordingly, this subsection starts with the description of a bilateral negotiation protocol. The protocol defines the tasks that two agents, represented generically by ag_1 and ag_2, can perform during the negotiation process.

This subsection also describes a multilateral negotiation protocol. The protocol defines the set of possible tasks that each agent $ag_i \in Ag$ can perform at each point of the negotiation process. A negotiation strategy specifies a particular task to perform from the set of possible tasks.

The Bilateral Negotiation Protocol. The process of negotiation starts with one agent, say ag_1, communicating a proposal $prop^{t1}_{1km}$ to the other agent ag_2. Next, ag_2 receives $prop^{t1}_{1km}$ and may decide either: (i) to accept $prop^{t1}_{1km}$, (ii) to reject $prop^{t1}_{1km}$, (iii) to make a critique $crit^{t2}_{2km}$ to $prop^{t1}_{1km}$, or (iv) to communicate a counterproposal $prop^{t2}_{2km}$. A *proposal* is a set of facts. A *critique* is a statement about issue priorities. A *counterproposal* is a proposal made in response to a previous proposal (see Subsection 4.3).

The process continues with ag_1 receiving the response of ag_2. Next, ag_1 checks whether an agreement was reached. If the proposal $prop^{t1}_{1km}$ was accepted, the process ends successfully. Otherwise, if ag_2 decided to reject $prop^{t1}_{1km}$ or to make a critique

$crit^{t2}_{2km}$, ag_1 can act either: (i) by communicating a new proposal $prop^{t3}_{1kn}$, or (ii) by sending an inform message acknowledging the receipt of ag_2's response. Otherwise, if ag_2 decided to communicate a counterproposal $prop^{t2}_{2km}$, ag_1 has the choice of either: (i) accept $prop^{t2}_{2km}$, (ii) reject $prop^{t2}_{2km}$, (iii) make a critique to $prop^{t2}_{2km}$, or (iv) communicate a new proposal $prop^{t3}_{1kn}$ (counterproposal).

The process of negotiation proceeds with ag_2 receiving the response of ag_1. The tasks just described are then repeated. The agents continue to negotiate until either: (i) they find an agreement, (ii) they reach a deadlock, or (iii) at least an agent decides to break off negotiation.

The Multilateral Negotiation Protocol. This protocol is similar to the previous protocol. The negotiation process starts with an agent, say ag_i, communicating a proposal $prop^{t1}_{ikm}$ to all the agents in A. Each agent $ag_j \in A$ receives $prop^{t1}_{ikm}$ and has the choice of either: (i) accept $prop^{t1}_{ikm}$, (ii) reject $prop^{t1}_{ikm}$ without making a critique, or (iii) reject $prop^{t1}_{ikm}$ and making a critique.

The process of negotiation proceeds with ag_i receiving the responses of all the agents in A. Next, ag_i checks whether an agreement was reached. If the proposal $prop^{t1}_{ikm}$ was accepted by all the agents in A, the negotiation process ends successfully. In this case, ag_i informs the agents in A that an agreement was reached. Otherwise, ag_i can act either: (i) by communicating a new proposal $prop^{t3}_{ikn}$, or (ii) by acknowledging the receipt of all the responses.

The process continues with the agents in A receiving the response of ag_i. If ag_i decides to communicate a new proposal $prop^{t3}_{ikn}$, each agent $ag_j \in A$ may again decide: (i) to accept $prop^{t3}_{ikn}$, or (ii) to reject $prop^{t3}_{ikn}$ without making a critique, or (iii) to reject $prop^{t3}_{ikn}$ and making a critique. If ag_i decides to acknowledge the receipt of the responses, the process continues to a new *round* in which another agent $ag_k \in Ag$ communicates a proposal to all the agents in $Ag - \{ag_k\}$. This is repeated for other agents in Ag.

The protocol does not make any assumption about who makes the first proposal, who is the second agent to make a proposal, and so on. Again, the agents negotiate until either: (i) they find an agreement, (ii) they reach a deadlock, or (iii) at least an agent decides to break off negotiation.

4.3. *The negotiation process* (*individual perspective*)

The individual model of the negotiation process specifies the tasks that each agent in Ag must perform in order to negotiate in an effective way. These tasks (or processes) are shown in Fig. 1 for the specific case of an agent $ag_i \in Ag$ that communicates a negotiation proposal. Let NP_{ik} represent ag_i's perspective of a negotiation problem and $NPstruct_{ik}$ be the structure of NP_{ik}. A formal description of the main processes follows.

Negotiation Proposal Generation. This process generates the set of negotiation proposals NPS_{ik} satisfying the requirements imposed by $NPstruct_{ik}$.

The generation of NPS_{ik} is performed through an iterative procedure involving three main tasks: (i) problem interpretation, (ii) proposal preparation, and (iii) proposal addition.

Problem interpretation consists of searching $NPstruct_{ik}$ for any solution sol_{ikm} of NP_{ik} and selecting the primitive plan templates of sol_{ikm}. More specifically, the search starts at the root node of $NPstruct_{ik}$, proceeds towards its leaves, and involves the arbitrary choice of exactly one plan template at each Or node of $NPstruct_{ik}$. This task is formalized by a function $interpret\_problem$ which takes $NPstruct_{ik}$ and NPS_{ik} as input and returns the primitive plan templates $PPT_{ikm} = \{pt_{ika}, \ldots, pt_{ikp}\}$ of sol_{ikm}.

Proposal preparation consists of determining a $negotiation\ proposal\ prop_{ikm} = \{f_{ika}, \ldots, f_{ikp}\}$, i.e., a set of facts corresponding to the headers of the primitive plan templates in PPT_{ikm}. This task is formalized by a function $prepare\_proposal$ which takes PPT_{ikm} as input and returns $prop_{ikm}$.

Proposal addition consists of adding a negotiation proposal $prop_{ikm}$ to the set NPS_{ik}. This task is formalized by a function $add\_proposal$ which takes NPS_{ik} and $prop_{ikm}$ as input and returns $NPS_{ik} + prop_{ikm}$.

It is worth to note that the preparation of a proposal $prop_{ikm}$ partitions the set F_{ik} of facts into: (i) subset $prop_{ikm} = \{f_{ika}, \ldots, f_{ikp}\}$, corresponding to the facts of a proposal, and (ii) subset $comp_{ikm} = \{f_{ikp+1}, \ldots, f_{ikz}\}$, called *complement* of $prop_{ikm}$, and corresponding to the remaining facts of F_{ik}.

The facts in $prop_{ikm}$ are fundamental for achieving the negotiation goal g_{ik}. They are the *inflexible facts* of negotiation, for proposal $prop_{ikm}$. The negotiation issues $Iprop_{ikm} = \{is_{ika}, \ldots, is_{ikp}\}$ associated with these facts are called *inflexible issues*. On the other hand, the facts in $comp_{ikm}$ are not important for achieving g_{ik}. They are the *flexible facts* of negotiation, for proposal $prop_{ikm}$. The issues $Icomp_{ikm} = \{is_{ikp+1}, \ldots, is_{ikz}\}$ associated with these facts are called *flexible* or *bargaining issues*.

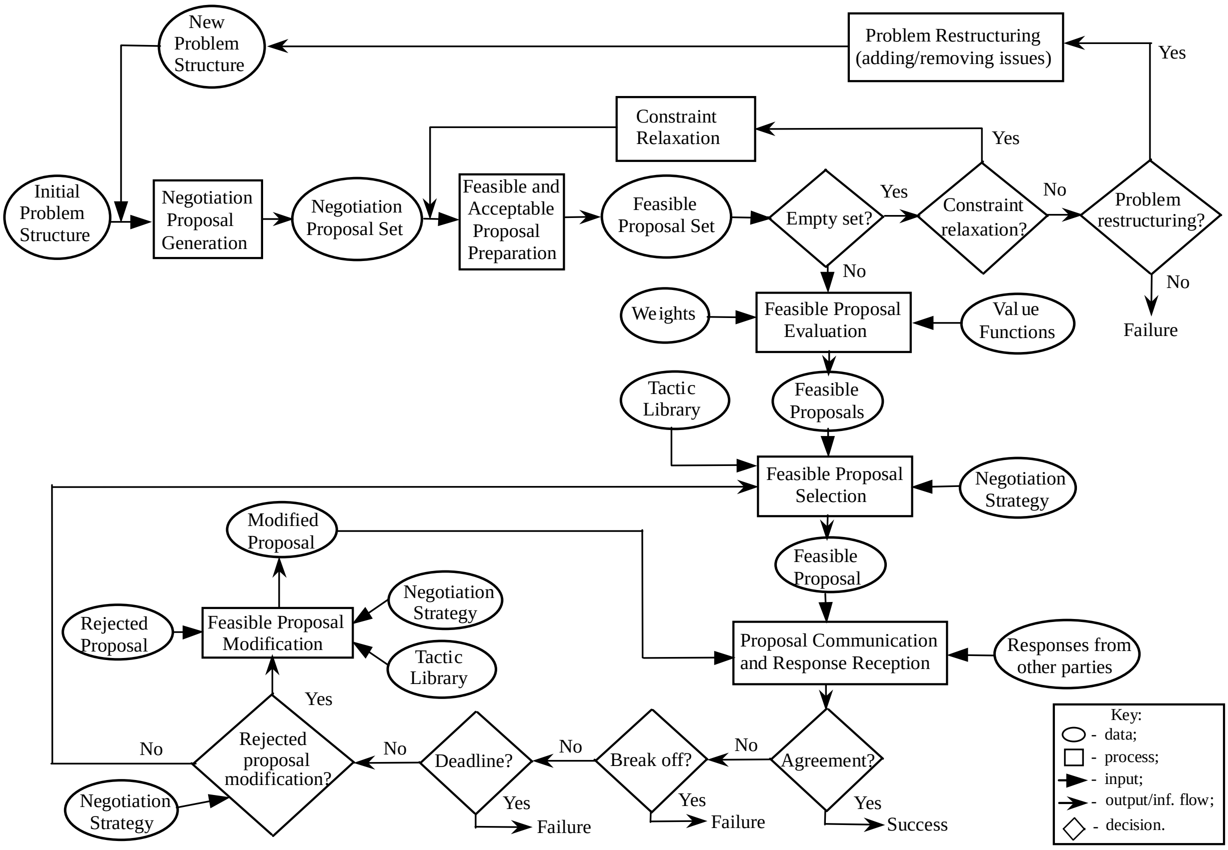

Fig. 1. The negotiation process (perspective of every agent that communicates a proposal).

Feasible and Acceptable Proposal Preparation. This process generates the set of feasible proposals $FPS_{ik}, FPS_{ik} \subseteq NPS_{ik}$, and the set of acceptable proposals $APS_{ik}, APSik \subseteq FPS_{ik}$.

Let $prop_{ikm} = \{f_{ika}, \ldots, f_{ikp}\}$ be a negotiation proposal. Let $Iprop_{ikm} = \{is_{ika}, \ldots, is_{ikp}\}$ be the set of issues associated with the facts in $prop_{ikm}$. Let $HCprop_{ikm} = \{hc_{ika}, \ldots, hc_{ikp}\}$ and $SCprop_{ikm} = \{sc_{ika}, \ldots, sc_{ikp}\}$ be the sets of hard and soft constraints for issues in $Iprop_{ikm}$, respectively. A negotiation proposal $prop_{ikm} \in NPS_{ik}$ is *feasible* if the issues in $Iprop_{ikm}$ satisfy the set $HCprop_{ikm}$ of hard constraints. A feasible proposal $prop_{ikm}$ is *acceptable* if the issues in $Iprop_{ikm}$ satisfy the set $SCprop_{ikm}$ of soft constraints.

The preparation of feasible proposals is formalized by a function $prepare\_feasible\_proposals$ which takes NPS_{ik} as input and returns FPS_{ik}. Similarly, the preparation of acceptable proposals is formalized by a function $prepare\_acceptable\_proposals$ which takes FPS_{ik} as input and returns APS_{ik}.

Feasible Proposal Evaluation. This process computes a score for each proposal in FPS_{ik} and orders the feasible proposals in a descending order of preference.

Let $prop_{ikm} = \{f_{ika}, \ldots, f_{ikp}\}$ be a feasible proposal. Let $Wprop_{ikm} = \{w_{ika}, \ldots, w_{ikp}\}$ be the set of weights of the issues in $Iprop_{ikm}$. Let $Cprop_{ikm} = (v[is_{ika}], \ldots, v[is_{ikp}])$ be the values of the issues in $Iprop_{ikm}$ ($Cprop_{ikm}$ is called a *contract*). The score of $prop_{ikm}$ is computed using the additive model [37]. For each issue $is_{ikj} \in Iprop_{ikm}, a \leqslant j \leqslant p$, let V_{ikj} be a function that gives the score ag_i assigns to a value $v[is_{ikj}]$ of is_{ikj} (V_{ikj} is called a *value function*). The score for contract $Cprop_{ikm}$ is given by the following expression:

$$V(Cprop_{ikm}) = \sum_{j=a}^{p} w_{ikj} V_{ikj}(v[is_{ikj}])$$

The proposal $prop_{ikm}$ is identified with contract $Cprop_{ikm}$ and both have the same score.

This process is formalized by a function $evaluate\_feasible\_proposals$. Let $W_{ik} = \{w_{ika}, \ldots, w_{ikz}\}$

and $VF_{ik} = \{V_{ika}, \ldots, V_{ikz}\}$ be the set of weights and value functions for the issues in I_{ik}, respectively. The function $evaluate\_feasible\_proposals$ takes FPS_{ik}, W_{ik} and VF_{ik} as input, computes a score $Vprop_{ikm} \in R$ for each feasible proposal $prop_{ikm} \in FPS_{ik}$, and returns the ordered set FPS_{ik}.

Feasible Proposal Selection. This process selects a feasible proposal from FPS_{ik}.

The process is formalized by a function $select\_feasible\_proposal$ which takes the set FPS_{ik}, the set APS_{ik}, the negotiation strategy str_{ik} and the library of tactics TL_i as input, and returns a proposal $prop^t_{ikj} \in FPS_{ik}$, where $t \in T$ denotes a generic instant of the negotiation process. The negotiation strategy str_{ik} dictates a specific tactic $tact_{ik} \in TL_i$ to use. The tactic $tact_{ik}$ specifies a particular proposal.

As stated in the previous subsection, the proposal that ag_i submits at the beginning of negotiation is denoted by $prop^{t1}_{ikm}$ and communicated to all the agents in A. If $prop^{t1}_{ikm}$ is not accepted by at least one agent in A, the agent ag_i may decide either: (i) to communicate a new proposal, or (ii) to acknowledge the receipt of all the responses. The new proposal can then be obtained either: (i) by selecting a new proposal $prop^{t3}_{ikn}$ from FPS_{ik}, or (ii) by modifying $prop^{t1}_{ikm}$.

The negotiation process continues with the agents exchanging more proposals. The proposal that ag_i submits at an instant t_n of the negotiation process is denoted by $prop^{tn}_{ikn}$.

Feasible Proposal Modification. This process computes a new proposal $prop^{tn+2}_{ikn}$ from a rejected proposal $prop^{tn}_{ikn}$.

The process is formalized by a function $modify\_rejected\_proposal$ which takes $prop^{tn}_{ikn}$, the negotiation strategy str_{ik} and the library of tactics TL_i as input and returns a new proposal $prop^{tn+2}_{ikn}$. The strategy str_{ik} defines one or two tactics to use. The tactics modify $prop^{tn}_{ikn}$ to make it more acceptable. The modification of $prop^{tn}_{ikn}$ can be done either: (i) by making a concession, or (ii) without making a concession.

4.4. Negotiation strategies

This subsection describes and formalizes two classes of strategies, called concession and problem solving strategies. The strategies are based on human negotiation procedures (see, for example [3,12,23]).

Concession strategies are functions that define the opening negotiation and concession tactics. The following three sub-classes of strategies are often used in real-world negotiations:

1. *starting high and conceding slowly* – these strategies model an optimistic opening attitude and successive small concessions;
2. *starting reasonable and conceding moderately* – these strategies model a realistic opening attitude and successive moderate concessions;
3. *starting low and conceding rapidly* – these strategies model a pessimistic opening attitude and successive large concessions.

The starting high and conceding slowly strategies are formalized by analogous functions. For instance, a strategy $SHCS1$ is formalized by a function $shcs1\_strategy$ which takes the library TL_i as input and specifies a tactic $tact_{ik}$ of a particular class $class\_tact_{ik}$:

$$shcs1\_strategy(TL_i) = (class\_tact_{ik},\ tact_{ik})|$$

$$\texttt{if}: state = \text{“}initial\text{”}\ \texttt{then}:$$

$$class\_tact_{ik} = \text{“}opening\ negotiation\text{”} \wedge$$

$$tact_{ik} = \text{“}starting\_optimistic\text{”}$$

$$\texttt{else}:$$

$$class\_tact_{ik} = \text{“}const\ conc\ factor\text{”} \wedge$$

$$tact_{ik} = \text{“}tough\text{”}$$

where $state = $ “$initial$” represents the initial state of the negotiation process (the beginning of negotiation), $starting\_optimistic$ is an opening negotiation tactic and *tough* is a constant concession factor tactic (see Subsection 4.5). The strategies in the other two sub-classes are formalized by functions essentially identical to that of $shcs1\_strategy$. These functions are, therefore, omitted.

The following six sub-classes of concession strategies are also used in real-world negotiations:

1. *starting high and conceding rapidly;*
2. *starting high and conceding moderately;*
3. *starting reasonable and conceding rapidly;*
4. *starting reasonable and conceding slowly;*
5. *starting low and conceding moderately;*
6. *starting low and conceding slowly.*

These strategies are only used in specific negotiation situations. They are similar to the previous strategies and their description and formalization are omitted (see, however, Subsection 5.2).

Problem solving strategies are functions that define the opening negotiation, concession and compensation tactics. The following two sub-classes of strategies are extensively used in real-life negotiations:

1. *low priority concession making* – these strategies model a realistic opening attitude, large concessions on issues of low priority and small concessions on other issues;
2. *low priority concession making with compensation* – these strategies are similar to the previous strategies; however, concessions are interleaved with compensations.

The low priority concession making strategies partition the set I_{ik} of issues into: (i) subset I_{ik}^{+}, corresponding to higher priority issues, and (ii) subset I_{ik}^{-}, corresponding to the remaining issues. The strategies in this sub-class are also formalized by analogous functions. For instance, a strategy $LPCM1$ is formalized by a function $lpcm1\_strategy$ which takes the library TL_i and the set I_{ik} as input, and returns the tactics $tact_{ik}$ and $tact_{ik+1}$ of classes $class\_tact_{ik}$ and $class\_tact_{ik+1}$, respectively:

$$lpcm1\_strategy(TL_i, I_{ik}) = (class\_tact_{ik},$$
$$tact_{ik}, I_{ik}^{+}, class\_tact_{ik+1}, tact_{ik+1}, I_{ik}^{-})|$$
$$\texttt{if}: state = \text{“}initial\text{”}\ \texttt{then}:$$
$$class\_tact_{ik} = \text{“}opening\ negotiation\text{”}\wedge$$
$$tact_{ik} = \text{“}starting\_realistic\text{”}\wedge$$
$$class\_tact_{ik+1} = \text{“}nil\text{”} \wedge tact_{ik+1} = \text{“}nil\text{”}$$
$$\texttt{else}: I_{ik} = I_{ik}^{+} + I_{ik}^{-}\wedge$$
$$class\_tact_{ik} = \text{“}const\ conc\ factor\text{”}\wedge$$
$$\forall it_{ikj} \in I_{ik}^{+},\ tact_{ik} = \text{“}tough\text{”}\wedge$$
$$class\_tact_{ik+1} = \text{“}const\ conc\ factor\text{”}\wedge$$
$$\forall it_{ikj} \in I_{ik}^{-},\ tact_{ik+1} = \text{“}soft\text{”}$$

where $state$ = “$initial$”, $starting\_optimistic$ and $tough$ have the meaning just specified, $starting\_realistic$ is an opening negotiation tactic and $soft$ is a constant concession factor tactic (see Subsection 4.5).

The formalization of the strategies in the other sub-class is essentially identical to that of $LPCM1$ and is omitted.

4.5. Negotiation tactics

This section describes and formalizes two classes of tactics, called opening negotiation and concession tactics. The tactics are also based on typical human negotiation procedures (see, for example [3,22,34]).

Opening negotiation tactics are functions that specify the proposal to submit at the beginning of negotiation.

Let $FPS_{ik} = \{prop_{ik1}, prop_{ik2}, \ldots, prop_{ikn}\}$ and $APS_{ik} = \{prop_{ik1}, prop_{ik2}, \ldots, prop_{ikh}\}$, $APS_{ik} \subseteq FPS_{ik}$, be the sets of feasible and acceptable proposals of ag_i, respectively. These sets are ordered in a descending order of preference. Let $prop_{ikh}$ be the acceptable proposal with the lowest score $Vprop_{ikh}$. Let Asp_{ikh} be the set of initial aspirations of ag_i for issues associated with the facts in $prop_{ikh}$. Let $Dif_{ikh} = |Vprop_{ikh} - VAsp_{ikh}|$, where $VAsp_{ikh}$ is the score of Asp_{ikh}.

Similarly, let $NAPS_{ik} = \{prop_{ikh+1}, \ldots, prop_{ikn}\}$, $NAPS_{ik} = FPS_{ik} - APS_{ik}$. Let $prop_{ikh+1}$ be the proposal of $NAPS_{ik}$ with the highest score $Vprop_{ikh+1}$. Let Asp_{ikh+1} be the set of initial aspirations of ag_i for issues associated with the facts in $prop_{ikh+1}$. Let $Dif_{ikh+1} = |Vprop_{ikh+1} - VAsp_{ikh+1}|$, where $Vasp_{ikh+1}$ is the score of Asp_{ikh+1}.

The following three tactics are used in many negotiation situations (for clarity, we omit the representation of time):

1. *starting optimistic* – specifies the proposal $prop_{ik1}$ with the highest score;
2. *starting realistic* – specifies either: (i) the proposal $prop_{ikh}$, if $Dif_{ikh} \leqslant Dif_{ikh+1}$, or (ii) the proposal $prop_{ikh+1}$, if $Dif_{ikh} > Dif_{ikh+1}$;
3. *starting pessimistic* – specifies the proposal $prop_{ikn}$ with the lowest score.

These tactics are formalized by similar functions. For instance, the tactic starting optimistic is formalized by the following function:

$$starting\_optimistic(FPS_{ik}) = prop_{ik1}|$$
$$\forall prop_{ikj} \in FPS_{ik}, Vprop_{ik1} \geqslant Vprop_{ikj}$$

The definition of the functions for the tactics starting realistic and starting pessimistic is essentially identical to that of $starting\_optimistic$ and is omitted.

Concession tactics are functions that compute new values for each issue at stake in negotiation. They model the concessions to be made on every issue at each point of the negotiation process.

Let I_{ik} be the negotiating agenda of ag_i. A *concession* on an issue $is_{ikl} \in I_{ik}$ is a change in the value of is_{ikl} that reduces the level of benefit sought. The *factor of concession* $Fc \in [0, 1]$ is a real number that defines the magnitude of every concession on is_{ikl}. We consider the following sub-classes of concession tactics:

1. *constant concession factor tactics* – model Fc as a constant;
2. *total concession dependent tactics* – model Fc as a function of the total concession made on is_{ikl}.

In each sub-class, we consider the following five tactics:

1. *stalemate* – models a null concession on is_{ikl};
2. *tough* – models a small concession on is_{ikl};
3. *moderate* – models a moderate concession on is_{ikl};
4. *soft* – models a large concession on is_{ikl};
5. *compromise* – models a complete concession on is_{ikl};

These tactics are often used by human negotiators.

Let $prop^{t1}_{ikm}$ be the proposal submitted by ag_i at the beginning of negotiation. Let $v[is_{ikl}]^{t1}$ be the value of is_{ikl} offered in $prop^{t1}_{ikm}$. Let V_{ikl} be a value function for is_{ikl} (this function is either monotonically increasing or monotonically decreasing). Let $V_{ikl}(v[is_{ikl}]^{t1})$ be the score of $v[is_{ikl}]^{t1}$.

Similarly, let $prop^{tn}_{ikn}$ be the proposal submitted by ag_i at an instant t_n of the negotiation process. Let $v[is_{ikl}]^{tn}$ be the value of is_{ikl} offered in $prop^{tn}_{ikn}$ and $V_{ikl}(v[is_{ikl}]^{tn})$ the score of $v[is_{ikl}]^{tn}$.

The *total concession* $tconc^{tn}_{ikl}$ made by ag_i on is_{ikl} at t_n is defined as follows:

$$tconc^{tn}_{ikl} = |v[is_{ikl}]^{t1} - v[is_{ikl}]^{tn}|$$

The *constant concession factor tactics* are formalized by a function $const\_factor\_tact$ which takes a value $v[is_{ikl}]^{tn}$ of is_{ikl}, the limit lim_{ikl} for is_{ikl} and two constants w and cte as input, and returns a new value $v[is_{ikl}]^{tn+2}$ for is_{ikl}:

$$\begin{aligned}&const\_factor\_tact(v[is_{ikl}]^{tn}, lim_{ikl}, w, cte)\\&\quad = v[is_{ikl}]^{tn+2} |\\&\quad v[is_{ikl}]^{tn+2} = v[is_{ikl}]^{tn} +\\&\quad (-1)^w Fc|lim_{ikl} - v[is_{ikl}]^{tn}| \wedge Fc = \text{“}cte\text{”}\end{aligned}$$

where $w = 0$ if V_{ikj} is monotonically decreasing or $w = 1$ if V_{ikj} is monotonically increasing. The five tactics are formalized by considering different values for Fc in the range $[0, 1]$.

The *total concession dependent tactics* are formalized by a function $total\_conc\_depd\_tact$ which takes a value $v[is_{ikl}]^{tn}$ of is_{ikl}, the limit lim_{ikl} for is_{ikl}, the total concession $tconc^{tn}_{ikl}$, the initial value $v[is_{ikl}]^{t1}$ of is_{ikl}, and two constants w and cte as input, and returns a new value $v[is_{ikl}]^{tn+2}$ for is_{ikl}:

$$\begin{aligned}&total\_conc\_depd\_tact(v[is_{ikl}]^{tn}, lim_{ikl}, tconc^{tn}_{ikl},\\&\quad v[is_{ikl}]^{t1}, w, cte) = v[is_{ikl}]^{tn+2} |\\&\quad v[is_{ikl}]^{tn+2} = v[is_{ikl}]^{tn} +\\&\quad (-1)^w Fc|lim_{ikl} - v[is_{ikl}]^{tn}| \wedge\\&\quad Fc = 1 - \lambda \left| \frac{tconc^{tn}_{ikl}}{lim_{ikl} - v[is_{ikl}]^{t1}} \right| \wedge\\&\quad \lambda = \text{“}cte\text{”}\end{aligned}$$

where $w = 0$ if V_{ikj} is monotonically decreasing or $w = 1$ if V_{ikj} is monotonically increasing, and $\lambda \in R^+$. The five tactics in this sub-class are formalized by considering different values for λ.

5. Experimental evaluation of the negotiation model

Experimentation mandates simplification [13]. Accordingly, the negotiation model is evaluated by performing a number of inter-related experiments. Each experiment empirically evaluates representative components of the model and lays the foundation for subsequent experimental work.

This section presents a detailed description of an experiment aiming at:

1. assessing the feasibility of building autonomous negotiating agents equipped with a simplified version of the negotiation model;
2. investigating the integration of planning and negotiation;
3. investigating the behavior of concession strategies and their associated opening negotiation and concession tactics; empirically evaluating these strategies and tactics by confirming a number of well-documented conclusions about human negotiation.

5.1. Empirical research on human negotiation

Much of the research on human negotiation concerns the effect of demand level and concession rate on the outcome of negotiation. A negotiator's demand level is the level of benefit to the self associated with the current offer. Concession rate is the speed at which demand level declines over time [34]. Most studies consist of laboratory experiments on two-party, single-issue negotiation. These studies support the following two conclusions [3,12,34]:

1. higher initial demands and slower concessions make agreement less likely and less rapidly reached;
2. lower initial demands and faster concessions produce smaller outcomes for the party employing them and larger outcomes for the other party, if agreement is reached.

These two conclusions imply a third, that there is an inverted U-shaped relationship between level of demand and the negotiation outcome:

1. negotiators who start with high demands and concede slowly often fail to reach agreement, which usually leads to inferior outcomes; those who start with low demands and concede rapidly usually reach agreement on the other party's terms, also yielding inferior outcomes; those between these extremes ordinarily achieve better outcomes.

The present study seeks to replicate these conclusions.

5.2. *The experimental system*

The experimental system consists of two autonomous agents and a simulated environment. Let $Ag = \{ag_s, ag_b\}$ be the set of agents. The agent ag_s plays the role of a seller (or a producer) and the agent ag_b the role of a buyer (or a customer). The agents negotiate the price of a generic commodity denoted by $prod_X$. A description of the agents and the environment follows.

Autonomous Negotiating Agents. Every agent $ag_i \in Ag$ is equipped with the model of individual behavior described in Subsection 3.1 and has a library CL_i of conflict detection axioms. We consider the following (for simplicity and clarity, we drop the subscripts k and j):

– the set G_i contains one goal – the agent ag_s has the goal g_s of selling $prod_X$ and the agent ag_b has the goal g_b of buying $prod_X$;
– the library PL_i contains five plan templates:
  (i) a plan template pt_{i1} representing a procedure for determining a price for $prod_X$;
  (ii) a plan template pt_{i2} for computing a perceived market value pmv_i for $prod_X$;
  (iii) three alternative plan templates pt_{i3}, pt_{i4}, and pt_{i5} for calculating a price pr_i for $prod_X$; each alternative plan template calculates pr_i from pmv_i, more specifically, by adding or subtracting a specific percentage of pmv_i to pmv_i;
– the intention structure IS_i contains one plan – the agent ag_s generates and adopts a plan p_s for achieving g_s and the agent ag_b generates and adopts a plan p_b for achieving g_b;
– the library CL_i contains the following axiom:

$$price(prod_X, pr_s) \& price(prod_X, pr_b) \& \neg(pr_s = pr_b) \rightarrow false$$

  where $price(prod_X, pr_s)$ and $price(prod_X, pr_b)$ represent the intentions of ag_s and ag_b to propose the prices pr_s and pr_b for $prod_X$, respectively.

Every agent ag_i is equipped with a simplified version of the negotiation model. The process of preparing and planning for negotiation involves the tasks specified in Subsection 4.1, except "negotiation strategy selection". This task is performed directly by the experimenter. The negotiation process of ag_i involves the five tasks specified in Subsection 4.3. We consider the following:

– the negotiating agenda I_i contains one issue is_i, namely the price of $prod_X$; the price ranges from $min_i = 0$ to $max_i = 1000$ currency units; the possible values of price are public information;
– the limit lim_i and the initial aspiration asp_i^{t1} for is_i are computed from pmv_i; the price specified in the proposal to submit at the beginning of negotiation is also computed from pmv_i;
– the agents are allowed to propose only strictly monotonically – the buyer's offers increase monotonically and the seller's offers decrease monotonically;
– the acceptability of a proposal is determined by a *negotiation threshold* – ag_i accepts a proposal $prop_j^{tn+1}$ when the difference between the price specified in $prop_j^{tn+1}$ and the price specified in the proposal $prop_i^{tn+2}$ that ag_j is ready to send is lower than or equal to the negotiation threshold of ag_i; the negotiation threshold of each agent is private information;
– the agents are allowed to exchange only a maximum number of proposals, denoted by max_{prop} – failure to reach agreement after max_{prop} proposals results in a deadlock; the parameter max_{prop} is public information.

The concession strategies and the associated opening negotiation and concession tactics of each agent ag_i are shown in Tables 1 and 2. In particular, Table 1 presents the six strategies used by both the seller and the buyer

Table 1
Negotiation strategies and tactics (for seller and buyer)

Agent	Strategy family	Strategy key	Opening negotiation tactic	Concession tactic family	Concession tactic
Seller and Buyer	Starting high and	SHCS1	Starting optimistic	Constant Concession Factor	Tough
	conceding slowly	SHCS2	Starting optimistic	Total Concession Dependent	Tough
	Starting reasonable and	SRCM1	Starting realistic	Constant Concession Factor	Moderate
	conceding moderately	SRCM2	Starting realistic	Total Concession Dependent	Moderate
	Starting low and	SLCR1	Starting pessimistic	Constant Concession Factor	Soft
	conceding rapidly	SLCR2	Starting pessimistic	Total Concession Dependent	Soft

Table 2
Negotiation strategies and tactics (only for buyer)

Agent	Strategy family	Strategy key	Opening negotiation tactic	Concession tactic family	Concession tactic
Buyer	Starting high and	SHCR1	Starting optimistic	Constant Concession Factor	Soft
	conceding rapidly	SHCR2	Starting optimistic	Total Concession Dependent	Soft
	Starting high and	SHCM1	Starting optimistic	Constant Concession Factor	Moderate
	conceding moderately	SHCM2	Starting optimistic	Total Concession Dependent	Moderate
	Starting reasonable and	SRCR1	Starting realistic	Constant Concession Factor	Soft
	conceding rapidly	SRCR2	Starting realistic	Total Concession Dependent	Soft
	Starting reasonable and	SRCS1	Starting realistic	Constant Concession Factor	Tough
	conceding slowly	SRCS2	Starting realistic	Total Concession Dependent	Tough
	Starting low and	SLCM1	Starting pessimistic	Constant Concession Factor	Moderate
	conceding moderately	SLCM2	Starting pessimistic	Total Concession Dependent	Moderate
	Starting low and	SLCS1	Starting pessimistic	Constant Concession Factor	Tough
	conceding slowly	SLCS2	Starting pessimistic	Total Concession Dependent	Tough

and Table 2 shows the twelve strategies used only by the buyer.

The constant concession factor tactics are applicable after the submission of the first proposal. However, the total concession dependent tactics are only applicable after the submission of the second proposal. Therefore, we consider the following:

– the agents compute the price to offer in the second proposal using a constant concession factor tactic. The price to offer in the third and subsequent proposals is computed accordingly to a negotiation strategy, i.e., using either a constant concession factor tactic or a total dependent concession tactic.

The Environment. The environment contains information about prior negotiations and market characteristics. This information is grouped into a single parameter called *base fair market value* and denoted by $bfmv_X$. We consider the following: (i) the value of $bfmv_X$ is public information, and (ii) $bfmv_X$ does not change throughout negotiation.

System Operation. The system operates in a simple and intuitive way. First, ag_s generates the plan p_s for achieving the goal g_s of selling $prod_X$. The plan p_s has a hierarchical structure that is embedded in the library PL_s. The perceived market value pmv_s is computed by randomly choosing a value within a specified percentage of the base $bfmv_X$. The price pr_s is set to pmv_s plus a percentage of pmv_s.

Next, ag_s writes the price pr_s of $prod_X$ to a public file. This procedure simulates the real-world procedure of advertising in appropriate places the desire to sell a product by a specific price.

Following this, ag_b generates a plan p_b for achieving the goal g_b of buying $prod_X$. The plan p_b is similar to plan p_s. Next, ag_b reads the price pr_s from the public file. This procedure simulates the real-world procedure of acquiring relevant information about a desired product. The agent ag_b then detects a conflict of interests $Conf$. The conflict arises because ag_b intends to buy $prod_X$ by pr_b, ag_s intends to sell $prod_X$ by pr_s, and $pr_b \neq pr_s$. Next, ag_b informs ag_s about the existence of the conflict. This is done by writing the conflict identifier $Conf$ to the public file. Next, ag_s is made aware of the conflict by reading $Conf$ from the public file.

The conflict is the driving force of negotiation. Therefore, the agents ag_s and ag_b start to negotiate a mutually acceptable agreement.

5.3. Experimental hypotheses

The experimental hypotheses postulate the replication of the conclusions presented in Subsection 5.1 and are stated as follows:

Hypothesis 1: The strategies SRCM1 and SRCM2 lead, on average, to higher payoffs than the strategies SHCS1 or SHCS2 and the strategies SLCR1 or SLCR2;

Hypothesis 2: The strategies SHCS1 and SHCS2 lead, on average, to slower agreements than the strategies SRCM1 or SRCM2 and the strategies SLCR1 or SLCR2;

Hypothesis 3: The strategies SHCS1 and SHCS2 lead, on average, to fewer agreements than the strategies SRCM1 or SRCM2 and the strategies SLCR1 or SLCR2.

5.4. The experimental method

The experimental method is controlled experimentation (see, for example [1,4]). A description of the experimental parameters, the independent variable, the dependent variables, and the experimental procedure follows.

Experimental Parameters. The agents and the environment have a built-in set of parameters that govern their behavior and facilitate experimentation. The relevant parameters for the experiment and their values are shown in Table 3. Most values are based on data and results of case studies published in the negotiation literature (e.g. [12,23,37]).

The base fair market value is set to 500 currency units. The perceived market value is generated by *randomly* choosing a value within a specified percentage of the base. This percentage is set to 10%. This models a system in which the market value is determined subjectively.

The limit and the initial level of aspiration are then computed from the perceived market value. The prices specified in a high, a moderate and a low initial offer are also computed from the perceived market value.

The Independent Variable. The independent variable is the preprogrammed strategy of the seller. This variable has six levels, namely the six strategies presented in Table 1. The value of this variable is under the control of the experimenter.

Dependent Variables. The dependent variables are the payoff that accrues to the seller, the time spend in negotiation and the outcome of negotiation. The values of these variables are not under the control of the experimenter. They are observed by the experimenter as measurements.

The first dependent variable is the payoff that accrues to the seller. The seller's payoff is a dependent variable because a major purpose of the research consists of examining the effect of concession strategies on the bargainer who uses these strategies (the seller in this study), and not on his opponent. Consider that ag_s and ag_b agree on a price pr. The payoff Vpr_s of ag_s for pr is given by the following linear function:

$$Vpr_s = pr - lim_s$$

where lim_s is the limit of ag_s for the price. If no agreement is reached in a particular negotiation, then the value of Vpr_s is set to zero.

The second dependent variable is the time spent in negotiation. This variable is measured in terms of the total number of offers exchanged by the agents until either they found an agreement or reach a deadlock. If no deal is made in a particular negotiation, then this variable is set to max_{prop}.

The last dependent variable is the outcome of negotiation (agreement or deadlock). This variable is used to compute the percentage of deals made in a number of negotiations.

The Experimental Procedure. The experiment involves six groups of trials. Each *group* corresponds to a level of the independent variable. A *trial* is a single run of the experimental system and involves a bargaining session. Trials of the same group will, in general, differ from one another, because the results of the system depend stochastically on the parameter settings, as stated above. The detailed experimental procedure is as follows:

1. for each group of trials:
   - 1.1 manipulate the independent variable (assign a strategy to the seller agent);
2. for each trial in each group:
   - 2.1 randomly determine the agent that starts the bidding process;
   - 2.2 randomly determine a strategy for the buyer agent;
   - 2.3 run the experimental system (allow the agents to negotiate using the specified strategies);
   - 2.4 measure the dependent variables;
3. for all trials of each group:
   - 3.1 compute averages on the measures taken in 2.4.

Table 3
Experimental parameter values

Experimental parameter	Value
Base fair market value	500 (currency unit)
Percentage for computing a perceived market value	10%
Percentage for computing the limit	50%
Percentage for computing the initial level of aspiration	35%
Percentage for computing a high initial offer (optimistic opening attitude)	55%
Percentage for computing a moderate initial offer (realistic opening attitude)	35%
Percentage for computing a low initial offer (pessimistic opening attitude)	15%
Maximum number of proposals	10
Negotiation threshold	5 (currency unit)

5.5. *Experimental results*

The experiment was conducted on a personal computer using Visual C^{++}. For each of the 6 groups, we conducted 31 trials. A pretest was performed to establish how many trials were needed to obtain significant averages on the measures taken (using both the analysis of variance and the Scheffé's method [4]). The experimental results are shown in Table 4.

The main response measure was the payoff that accrued to the seller. It was predicted that the strategies SRCM1 and SRCM2 yielded superior outcomes. Table 4 reports all the payments received by the seller (including those corresponding to a zero payoff). These results indicate that the strategy SRCM1 resulted in significantly higher payoffs when compared to the payoffs resulting from the strategies SHCS1 and SLCR1 ($F = 8.984,\ p < 0.05$). The same is true for the strategies SRCM2, SHCS2 and SLCR2 ($F = 14.282,\ p < 0.05$). Hypothesis 1 is supported.

The number of proposals exchanged by the agents was also recorded. The prediction was that the tougher the seller, the higher would be the number of proposals the agents would exchange for an agreement to be reached. The results indicate that this prediction was confirmed. The strategy SHCS1 resulted in significantly more proposals than the strategies SRCM1 and SLCR1 ($F = 151.986,\ p < 0.005$). The same is true for the strategies SHCS2, SRCM2 and SLCR2 ($F = 134.178,\ p < 0.005$). Hypothesis 2 is also supported.

The last measure taken was the number of cases when agreement was reached. The prediction was that the tougher the seller, the higher would be the number of cases when no agreement was reached. The results show that this prediction was also confirmed. The strategies SHCS1 and SHCS2 led to fewer agreements. Hypothesis 3 is, therefore, supported.

6. A survey of existing negotiation models

Negotiation is a rich, multidisciplinary research area. Hence, our purpose in this section is not to provide a comprehensive overview, but rather to compare our model with other developed models.

Laasri et al. [21] present a generic model of negotiation. The model assumes that the agents pursue common goals and are cooperative.

Rosenschein and Zlotkin [39] use game theory to investigate the properties of negotiation protocols. Their work does not make the cooperating agent assumption. However, it embodies a number of limiting assumptions. In particular, it assumes that the agents have complete knowledge of the other agents' preferences.

Sycara [40] presents a negotiation model that can be employed by non-cooperative agents and supports problem restructuring. However, the model assumes the existence of a centralized mediator. Kraus et al. [19] extend the work of Sycara and present a logical model of the process of argumentation. Their work concentrates on developing a new logic, defining a number of arguments and implementing an automated negotiation agent. Therefore, no consideration was given to dynamically change negotiation proposals and to introduce new issues.

Faratin et al. [7] present a multi-party, multi-issue model of negotiation. The model is based on computationally tractable assumptions and empirically evaluated. However, no consideration was given to integrate the model with existing models of individual behavior.

We are interested in negotiation among both self-motivated and cooperative agents. Our negotiation model is generic and supports both dynamic constraint relaxation and problem restructuring. Our representation for negotiation problems is similar to decision trees and goal representation trees [11,16]. There are, however, important differences. Our approach does not require the quantitative measures typical of decision analysis. Also, our approach is based on plan templates

Table 4
Experimental results

Seller's strategy	Seller's payoff (mean)	Number of proposals (mean)	Percentage of agreements
SHCS1	121.806	9.483	32.258
SHCS2	122.935	9.516	32.258
SRCM1	243.258	6.709	93.548
SRCM2	259.225	6.419	100.000
SLCR1	157.193	4.193	100.000
SLCR2	141.193	4.645	100.000

and plan expansion, and not on production rules and forward or backward chaining.

Our negotiation model defines and formalizes a number of negotiation strategies and tactics. Our formulae for modeling concession tactics are similar to the formulae used by other researchers [7,17]. Again, there are important differences. Our formulae assure that the new value of an issue always ranges between the limit and the previous value of the issue. Also, our formulae are based on the total concession made by an agent on an issue, a criterion not used by other researchers. Finally, our formulae model a number of well-documented conclusions about the effect of demand level and concession rate on the outcome of negotiation.

7. Conclusion

This paper presented a computational negotiation model for autonomous agents. There are several features of the model that should be highlighted. First, the model is generic and can be used in a wide range of domains. Second, the model acknowledges the role of conflict as a driving force of negotiation. Third, the model accounts for a tight integration of the individual and social behavior of agents. In particular, the structure of a problem allows the direct integration of planning and negotiation. This structure also defines the set of negotiation issues. Fourth, the model supports problem restructuring. This feature assures a high degree of flexibility. More specifically, problem restructuring facilitates the removal of deadlocks and increases the parties' willingness to a compromise. Problem restructuring also allows the dynamic addition of negotiation issues. Finally, the negotiation strategies and tactics are motivated by human negotiation procedures.

This paper also described an experiment performed to evaluate empirically a number of representative components of the model. The experimental results showed that: (i) the strategies of the class starting reasonable and conceding moderately lead, on average, to superior outcomes, and (ii) the strategies of the class starting high and conceding slowly lead, on average, to fewer and slower agreements. The results confirmed a number of basic conclusions about human negotiation.

Our aim for the future is to continue the development of the negotiation model and to extend the experimental evaluation of the model. In particular, we intend to add a number of negotiation strategies and tactics and to consider problem restructuring. We also intend to perform an experiment to investigate the behavior of problem solving strategies and to evaluate the effectiveness of these strategies. In addition, we intend to perform a number of experiments to observe the differences between agents that dynamically change the representation of negotiation problems and agents that use fixed representations.

References

[1] E. Aronson, T. Wilson and M. Brewer, Experimentation in Social Psychology, in: *Handbook of Social Psychology,* (Vol. 1), D. Gilbert, S. Fiske and G. Lindzei, eds, McGraw-Hill, 1998, pp. 99–142.

[2] M. Bratman, *Intentions, Plans and Practical Reason*, Harvard University Press, Cambridge, MA, 1987.

[3] P. Carnevale and D. Pruitt, Negotiation and Mediation, in: *Annual Review of Psychology,* (Vol. 43), M. Rosenzweig and L. Porter, eds, Annual Reviews Inc., 1992, pp. 531–581.

[4] P. Cohen, *Empirical Methods for Artificial Intelligence,* MIT Press, Cambridge, MA, 1995.

[5] R. Conte and C. Castelfranchi, *Cognitive and Social Action,* UCL Press, London, 1995.

[6] K. Erol, *Hierarchical Task Network Planning: Formalization, Analysis, and Implementation,* Ph.D. Dissertation, University of Maryland, USA, 1995.

[7] P. Faratin, C. Sierra and N. Jennings, Negotiation Decision Functions for Autonomous Agents, *Robotics and Autonomous Systems* **24**(3–4) (1998), 159–182.

[8] I. Ferguson, *Touring Machines: An architecture for Dynamic, Rational Agents,* Ph.D. Dissertation, University of Cambridge, UK, 1992.

[9] J. Galliers, *A Theoretical Framework for Computer Models of Cooperative Dialogue Acknowledging Multi-Agent Conflict,* Ph.D. Dissertation, University of Cambridge, UK, 1989.

[10] M. Georgeff, B. Pell, M. Pollack, M. Tambe and M. Wooldridge, The Belief-Desire-Intention Model of Agency, in: *Intelligent Agents V,* Springer Verlag, Berlin, 1999 (LNAI 1365).

[11] P. Goodwin and G. Wright, *Decision Analysis for Management Judgement,* John Wiley&Sons, 1991.

[12] W. Hamner, Effects of Bargaining Strategy and Pressure to Reach Agreement in a Stalemated Negotiation, *Journal of Personality and Social Psychology* **30**(4) (1974), 458–467.

[13] S. Hanks, M. Pollack and P. Cohen, Benchmarks, Testbeds, Controlled Experimentation, and the Design of Agent Architectures, *AI magazine* **14**(4) (1993), 17–42.

[14] B. Hayes-Roth, On Building Integrated Cognitive Agents: A Review of Allen Newell's Unified Theories of Cognition, *Artificial Intelligence* **59**(1–2) (1993), 329–341.

[15] J. Horty and M. Pollack, Option Evaluation in Context, in: *Theoretical Aspects of Rationality and Knowledge,* 1998.

[16] G. Kersten, W. Michalowski, S. Szpakowicz and Z. Koperczak, Restruturable Representations of Negotiation, *Management Science* **37**(10) (1991), 1269–1290.

[17] Z. Koperczak, S. Matwin and S. Szpakowicz, Modelling Negotiation Strategies with Two Interacting Expert Systems, *Control and Cybernetics* **21**(1) (1992), 105–130.

[18] S. Kraus, J. Wilkenfeld and G. Zlotkin, Multiagent Negotiation under Time Constraints, *Artificial Intelligence* **75** (1995), 297–345.

[19] S. Kraus, K. Sycara and A. Evenchick, Reaching Agreements Through Argumentation: A Logical Model and Implementation, *Artificial Intelligence* **104**(1–2) (1998), 1–69.

[20] M. Klusch, *Intelligent Information Agents,* Springer-Verlag, Berlin, 1999.

[21] B. Laasri, H. Laasri, S. Lander and V. Lesser, A Generic Model for Intelligent Negotiating Agents, *Journal of Intelligent and Cooperative Information Systems* **1**(2) (1992), 291–318.

[22] R. Lewicki and J. Litterer, *Negotiation,* Irwin, Homewood, Illinois, 1985.

[23] R. Lewicki, D. Saunders and J. Minton, *Negotiation, Readings, Exercises, and Cases,* McGraw Hill-Irwin, Boston, 1999.

[24] F. Lopes, N. Mamede, H. Coelho and A.Q. Novais, A Negotiation Model for Intentional Agents, in: *Multi-Agent Systems in Production,* Elsevier Science, Amsterdam, 1999, pp. 211–216.

[25] F. Lopes, N. Mamede, A.Q. Novais and H. Coelho, Towards a Generic Negotiation Model for Intentional Agents, in: *Agent-Based Information Systems,* IEEE Computer Society Press, CA, 2000, pp. 433–439.

[26] F. Lopes, N. Mamede, A.Q. Novais and H. Coelho, Conflict Management and Negotiation Among Intentional Agents, in: *Agent-Based Simulation,* SCS-Europe, 2001, pp. 117–124.

[27] F. Lopes, N. Mamede, A.Q. Novais and H. Coelho, Negotiation Tactics for Autonomous Agents, in: *Internet Robots, Systems and Applications,* IEEE Computer Society Press, CA, 2001, pp. 708–714.

[28] F. Lopes, N. Mamede, A.Q. Novais and H. Coelho, Negotiation Among Autonomous Computational Agents, in: *Iberamia 2002 – IberoAmerican Conference on Artificial Intelligence,* Springer Verlag, Berlin, 2002 (LNAI 2527), pp. 556–565.

[29] D. McDermott and J. Hendler, *Artificial Intelligence,* 76 (special volume on planning and scheduling), 1995.

[30] J. Muller, *The Design of Intelligent Agents,* Springer-Verlag, Berlin, 1996 (LNAI 1177).

[31] M. Osborne and A. Rubinstein, *Bargaining and Markets,* Academic Press, San Diego, CA, 1990.

[32] S. Parsons, C. Sierra and N. Jennings, Agents that Reason and Negotiate by Arguing, *Journal of Logic and Computation* **8**(3) (1998), 261–292.

[33] M. Pollack, Plan Generation, Plan Management and the Design of Computational Agents, in: *ICMAS'98 -International Conference on Multi-Agent Systems,* 1998.

[34] D. Pruitt, *Negotiation Behavior,* Academic Press, London, 1981.

[35] D. Pruitt, Social Conflict, in: *The Handbook of Social Psychology,* (Vol. 2), D. Gilbert, S. Fiske and G. Lindzei, eds, McGraw-Hill, 1998, pp. 470–503.

[36] D. Pruitt and J. Rubin, *Social Conflict: Escalation, Stalemate and Settlement,* McGraw-Hill, New York, USA, 1986.

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

[38] A. Rao, Integrated Agent Architecture: Execution and Recognition of Mental-States, in: *Distributed Artificial Intelligence, Architecture and Modelling,* C. Zhang and D. Lukose, eds, Springer Verlag, Berlin, 1996, pp. 159–173.

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

[40] K. Sycara, Problem Restructuring in Negotiation, *Management Science* **37**(10) (1991), 1248–1268.

[41] F. von Martial, *Coordinating Plans of Autonomous Agents,* Springer-Verlag, Berlin, 1992.

[42] D. Zeng and K. Sycara, Bayesian Learning in Negotiation, *International Journal of Human-Computer Studies* **48** (1998), 125–141.