Proposal for a multi-agent model with a human-assisted adaptation and learning mechanism

RUDDY LELOUCHE*
Université Laval

STÉPHANE DOUBLAIT**
Bell Atlantic International

Abstract

Very few attempts have been made to address the problem of modeling and adapting complex mental attitudes in multi-agent environments. In this paper, we introduce the concept of scene static description and propose a categorization of actors' mental and physical attitudes, based on the concepts of role, status, and behavior. We also model background knowledge of a scene through spatial and temporal statements. This categorization allows an elegant modeling of actors' beliefs and bluff intentions (including recursive ones) directed towards other actors. Beliefs and bluff intentions are two critical characteristics of conflicting multi-agent environments. Based on the assumption that the use of abstraction facilitates and enhances reasoning, we introduce four multi-level hierarchies (descriptors, beliefs, goals, and plans) to host the knowledge elements previously categorized. We also introduce transitions between these hierarchies to create an analysis-resolution reasoning cycle, aimed at finding for a particular actor, the best possible sequence of behaviors to adapt to a constantly chnging environment. Adaptation may also occur through a human-assisted learning mechanism which we also define in this paper. This mechanism follows the learning-by-instruction paradigm, and is illustrated in the context of a simulation which involves a human tutor to interactively assist an automatic player at a four-player strategy game. This form of learning addresses the problems of theory over-specificity and over-generality, and handles multiple faults.

Propuesta de un modelo multi-agente con un mecanismo de adaptación y aprendizaje asistido por humanos

Resumen

Pocas veces se ha afrontado el problema de modelar y adaptar actitudes mentales complejas en entornos multi-agente. En este artículo, introducimos el concepto de descripción estática de escenas y proponemos una categorización de las actitudes mentales y físicas de los actores, basada en los conceptos de rol, estatus y conducta. También modelamos el conocimiento subyacente a una escena a partir de declaraciones espaciales y temporales. Esta categorización permite un modelaje elegante de las creencias y de las intenciones de engaño de los actores (incluidas las recursivas) dirigidas hacia otros actores. Las creencias e intenciones de engaño son rasgos críticos de los entornos multi-agente conflictivos. Basándonos en el supuesto de que el uso de abstracciones facilita y mejora el razonamiento, introdujimos cuatro jerarquías multi-nivel (descriptores, creencias, metas y planes) para organizar los elementos de conocimiento previamente categorizados. Además introdujimos transiciones entre estas jerarquías con el fin de crear un ciclo de razonamiento análisis-resolución, destinado a hallar para un actor particular, la mejor secuencia posible de conductas para adaptarse a un entorno en constante cambio. La adaptación también puede ocurrir en un mecanismo de aprendizaje asistido por humanos que definiremos en este artículo. Dicho mecanismo sigue el paradigma de aprendizaje-por-instrucción, y se ilustra en el contexto de una simulación que implica un tutor humano que asiste interactivamente a un jugador automático en un juego estratégico de cu atro-jugadores. Esta forma de aprendizaje afronta los problemas de hiper-especificidad e hiper-generalidad de la teoría, y maneja defectos múltiples.

Authors's Address: *Departament of Computer Science. Université Laval. Sainte-Foy, Québec G1K 7P4. Canadá. **Bell Atlantic International. 1310 North Courthouse Road, 5 th floor Arlington, VA 22209. USA.

 Cognitiva 1993, *5 (2)*, 245-269

1. INTRODUCTION

In real life, learning and adaptation are the result of a repeated cycle between physical observable facts and mental attitudes. The observable facts are the ones that are perceived by the person, or actor. They consist of background information and characteristics as well as behaviors of other actors. The mental attitudes are the results of this perception, and depending on the truthfulness of the corresponding information, we shall call them knowledge (true information), beliefs (information that may be true or false depending on the accuracy of the belief), and bluff (information that is false by definition).

Figure 1 shows how the cycle proceeds for the actor whom we are dealing with, and whom we shall call the reference actor. Starting from some piece of observable data, the reference actor interprets this data, together with the mental attitudes already present in his or her mind, and elaborates on them, in order to build new mental attitudes (learning) and to drop or modify (adaptation) existing ones. This reasoning is an analytical process, based on the interpretation of observable facts and behaviors and aimed at generating mental attitudes. Later, or at a different time, the existing actor's mental attitudes can be made explicit to the outside world, through the generation of externally observable behaviors. In the simplest case, such behaviors merely act upon the environment in order to modify it. In more complex and interesting cases, these behaviors are aimed at making something known to others. Then, depending on the purpose or goal to be attained by the reference actor, the purpose of such explicitation can be merely to communicate with others, or to fool them (if some type of bluff is needed). In all cases, the second half of the cycle is concerned with adaptation to the environment (in the broadest sense) and involves a synthetical type of reasoning, based on the explicitation of mental attitudes and aimed at generating behaviors and other observable facts.

FIGURE 1.
The learning and adaptation cycle

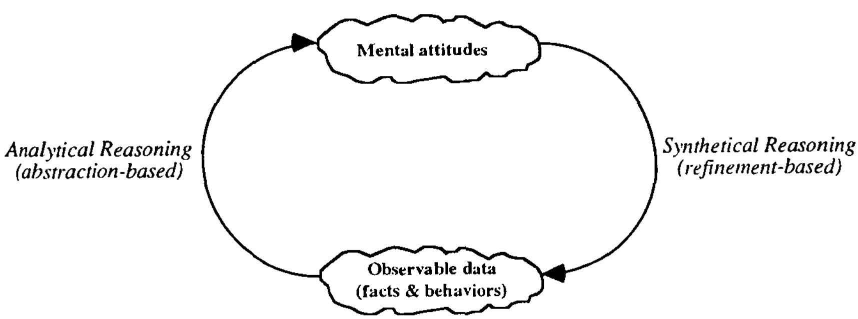

Section 2 of this paper introduces the pieces on which our model proposal is built, i.e. the classification of the various knowledge elements and structures which we believe should be present in a true multi-agent environment, and are necessary to detail and possibly implement the reasoning cycle shown in figure 1. Then section 3 concentrates on the processes involved in the two halves of this reasoning cycle (arrows of figure 1) and shows how abstraction levels can be used to adequately model them. Finally, in section 4, we attempt to show on a multi-agent simulation (a four-player strategy gama named Sigma-File) how these reasoning processes can be automated and further refined with the help of a human tutor through a learning-by-instruction mechanism.

2. TAXONOMY OF KNOWLEDGE ELEMENTS IN A MULTI-AGENT ENVIRONMENT

In this section we first explicit what we call knowledge, belief, and bluff, and how they are related to mental attitudes. We then will introduce in 2.2. the concept of scene static description, and summarize in the rest of the section the various types of knowledge elements that characterize complex multi-agent environments. A complete description of this approach, including formalization aspects, using first-order logic, can be found in (Lelouche & Doublait, 1992) (in that text, «multi-actor» is used equivalently with «multi-agent»).

2.1. Knowledge, beliefs, and bluff

We use the term «information» as a generic term, to designate some knowledge that can be true or false. Actually, this «information» can be divided into three categories: knowledge, belief and bluff. We shall use the term knowledge to refer to justified true information (as it is commonly defined, as for instance in (Fetzer, 1985), the term belief to refer to some information that may be true or false (whatever the cause), but that is considered as true by the owner of that information (a particular actor), and the term bluff to refer to false information intentionally introduced as true by and actor into some other actor's information base, or in some other actor's view of the public information base. That definition of bluff implies that the number of actors in the scene must be at least two, since bluff requires one actor to introduce it and at least one actor to believe it (as true or false): thus, a multi-agent environment is a necessary condition to bluff modeling.

Observable facts and behaviors are obviously a particular case of knowledge. However, mental attitudes may consist of any of the three types.

2.2. The concept of scene static description

In our model, we first consider the environment to be a colection of scene static descriptions described in (Lelouche, 1986), each description being made of a background (temporal an geographical context) and of some actors. We use the term static because a static description can be seen as a snapshot of a situation or scene, temporarily freezing the background and the actors' mental attitudes (knowledge, beliefs, and bluff intentions). Moving from one description to another is accomplished through the execution of actors' behaviors or planned sequences of behaviors that may change the background and the actors' mental attidues.

In our scene static description model, we distinguish various kinds of information, as illustrated in figure 2. In the remainder of this section, we describe the nature of each of these categories of information.

FIGURE 2.
Categories of scene information

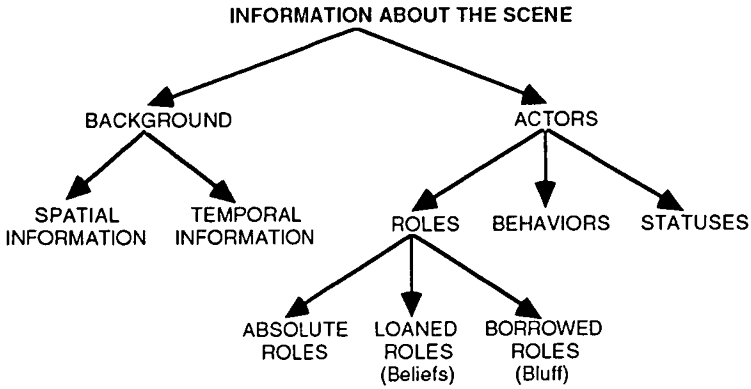

2.3. Background information

Background information is divided into two groups: spatial information and temporal information. Spatial information asserts facts about spatial properties of the scene itself such as «there is a table in the room», «it is raining outside», or «spy Ivanov is located in Tangier» (in the Sigma-File game). Note that in the last example, spies are not considered as actors, but as objects of the scene. Temporal information asserts facts about temporal properties of the scene itself such as «the time of day is morning» or «the current game turn number is 17».

Thus, the background is all information that is unrelated to actors: this is indeed its true and compte definition. But in this paper, we also assume that it is stable in time and available to every actor. These assumptions do not hold in general. For example, it might be raining now, but sunny later or tomorrow. Also I might not know the size of a room if I have never seen this room and nobody ever described it to me. Moreover, I may believe that there is a table in that room is somebody told me so, whether it is true or not. Therefore, the possibility of belief and even bluff about the back-ground do exist in real life.

However, in this paper, we mainly deal with information concerning actors, because it is the most interesting, but also the most difficult to adequately model. This is why, in order to simplify our discussion, we assume that all information concerning the background is observable and certain. Indeed, the possibility of belief and bluff about the background can be easily inferred, when it is adequate to do so, from our discussion limited to those cases where beliefs and bluff only deal with actors (who may be oneself or other actors).

2.4. Actors' statuses and behaviors

By contrast to the environment background, actors are capable to make decisions, and the execution of such decisions might lead to changing the environment. Information concerning actors is divided into three categories: statuses, behaviors and roles. Here, we discuss the first two categories, which are the simplest. We then discuss roles in section 2.5.

2.4.1. *Actors' statuses*

Every actor is defined by a set of statuses that currently or permanently define him, independently from the other actors, but possibly depending on the background. Examples of statuses are: «John is a male», «John has got the measles» (Health status), «John is in room-1» (geographical status), «John is a doctor» (social status), «player CIA has a current balance of $2500» (financial status in the Sigma-file game), or «player KGB is to play next» (game status).

2.4.2. *Actors' behaviors*

Behaviors reflect actions that an actor may perform to make the scene change. A behavior is a static piece of information in the sense that it corresponds to a physical act that might later be used to infer new mental actor's attitudes. Examples of behaviors are «Bob is sitting», «robert is carrying an umbrella», «John is kissing Mary», or «player CIA moves spy Ivanov from Tangier to Paris».

Like statuses, some behaviors necessarily imply another actor or other actors (e. g. to kiss or to teach), whereas other ones could take place in a single actor environment (position, clothing, movements, etc.). But the difference between a status and a behavior when others are concerned is that the execution of the behavior implies the physical presence of that or those specific actor(s), whereas the status only implies his/her existence.

2.5. Actors' roles

In our model, roles are the most complex category of information about actors and thus deserve to be addessed in more detail. A role defines possible relationships between an actor and one (or more) of his effective protagonist. Following (Lelouche, 1986), we distinguish three types or roles: absolute roles, loaned roles and borrowed roles.

2.5.1. *Absolute roles*

Absolute roles are the most obvious. They model various kinds of relationships between actors (e.g. social, psychological, relational, affective). Examples of absolute roles are «John loves Mary» (an affective role known as true by John), «player CIA cooperates with player IS» (a relational role known as true at least by player CIA), or «John is Mary's boss» (a social and relational role known by both John and Mary).

The main difference between an absolute role and a status is that the latter defines an actor independently from the other actors in the scene (even if the status implies the existence of other actors), whereas an absolute role defines him with respect to some other specific actor(s). Thus «John is a professor» is a social (permanent) status since it does not depend on the actor with whom John is in relation at a specific moment, but «John is Mary's professor» is an absolute role, specifically directed to one particular actor. However, the physical presence of Mary is not necessary for that role to hold, like a behavior would imply.

2.5.2. *Beliefs and loaned roles*

A belief denotes some information which is part of an actor's universe, whether this information is true or not, i.e. whether it actually holds, but which that actor holds as true. For example, I might believe that there is a chair in a room that I don't know (whether there is one or not), or that it will be raining tomorrow (whatever the weather). Such beliefs concern the environment.

Since a belief is uncertain by nature, it may be convenient to associate to it some level of confidence. The latter my be either quantitative (using numeric values similar to certainty factors (Lenat & Guha, 1990), or purely qualitative (using fuzzy adjectives). For instance, «Bill is John's professor» could be induced with maybe 60% of confidence, or expressed qualitatively, with a «fair» level of confidence.

However, there also exist beliefs about (other) actors. In order to introduce that particular notion, we define the concept of loaned roles. A role is said to be loaned by an actor X to an actor Y when it represents some (possibly uncertain) perception by X, of some role or status of Y. For instance, if John thinks that Mary loves him, «Mary loves John» denotes a role loaned by John to Mary, which may or may not be consistent with one of Mary's absolute roles.

Thus, beliefs can concern either the background or other actors. Since in this paper we are interested mainly in the latter, as already mentioned, we asume that all information concerning the environment is accessible to everyone (although this is not generally the case in real life), and therefore that there is no untrue belief about the environment.

2.5.3. *Bluff and borrowed roles*

Similarly, we saw that bluff also may concern the environment, but we discard this case here. So, in order to introduce the notion of bluff towards actors, we define the concept of borrowed roles. A role is said to be borrowed by an actor X towards an actor Y when it conveys information that in general is false but that X wants Y to believe is true. An actor's borrowed role might be (and often is) inconsistent with an absolute role. For instance, William may pretend to love Mary (a borrowed role) while in fact he loves only Jane (an absolute role).

2.5.4. *Role recursivity*

To model adequately and expressively the actors' loaned and borrowed roles, we adopt the recursive cognitive solipsism (a philosophical term introduced by Wilks (1985) in Artificial Intelligence for beliefs), which means that we are modeling in a recursive fashion some particular actor's roles (his beliefs but also his bluff intentions). We then allow an arbitrary nesting level of roles to represent complex expressions such as «John suspects that Mary wants William to believe that Jane is a doctor». Barnden (1989) mentions that researchers usually de-emphasize the problem of dealing in detail with nested attitudes (e.g. belief about a belief or a status) in favor of concentrating on the sub-issue of non-nested attitudes, and then extend their approach in a simple way to handle nested attitudes. However he also remarks that a wide variety of schemes for attitude representation are prone to a deep but subtle problem when they are applied to nested attitudes. We agree completely with Barnden on this point and already decided to handle nested attitudes directly in our formalism. Nevertheless, our main reason for doing so is that loaned and borrowed roles are inherently nested attitudes and can thus never be expressed in a non-nested way.

In fact, the deepest (terminal) nesting level of a recursive (i.e. loaned or borrowed) role is necessarily an absolute role, a status or a behavior, but it can never be a background piece of information. The reason is that loaned roles only model a subset of what and actor may believe, namely beliefs about other actors, and cannot model concerning the background, such as «John believes that there is a table in the room». Similarly, borrowed roles can model only bluff concerning other actors, and therefore cannot model bluff intentions such as«John wants William to believe that there is a table behind him» (a bluff concerning the background). The reason for this restriction lies in the very definition of loaned and borrowed roles: A role is loaned by an actor to some other actor or actors and a role is borrowed by an actor towards some other actor or actors. For instance, an actor cannot loan to one of his protagonists the role corresponding to the presence or absence of a table in a room (but he can loan him the role of being a doctor, i. e. a status). These results may seem confusing since we have a natural tendency to express loaned and borrowed roles with verb phrases like «... believes that...» or «... thinks that...» (for loaned roles), and «...wants...to believe that...» or «...pretends to...» (for borrowed roles). We shall keep such phrases for the sake of clarity, because using a verb phrase like «...loans to...the role of...», although more precise, would burden the natural language.

2.6. Plans as sequences of behaviors

In order to explicit some mental attitude, the reference actor must exhibit some kind of behavior(s), which may be more or less elaborate, but which must be visible, i.e. observable by other actors. We assume that the triggering mental attitude eventually leads to some goal, as we make this clearer in section 3. Such possible goals include, but are not limited to, the following: having something changed in the background, having some absolute role (eg. a relational or psychological role) demonstrated, or having some belief (whether true or false) transferred in someone else's state of mind. The way such a goal may be attained depends on its complexity, and the chances that it is actually attained also depend on whether it is true knowledge, belief, or bluff.

In any case, we assume that the means through which the expected goal is pursued is a behavior of the reference actor. Since that behavior may be more or less elaborate, we also assume that, in the most complex case, it consists of a structured sequence of elementary behaviors. Such a sequence may be very intricate. Indeed, inducing a false knowledge in somebody else's mind requires more time and a more complex sequence of behaviors than «simply» acting upon the «passive» background[1]. In general, such a sequence is assumed to be the result of some strategy, or tactics, that the reference actor is to follow. In the sequel, to use as neutral a term as possible and to follow a widely accepted terminology (although it is most often used in other domains), we refer to this process as a plan, and the activity taking place between the goal appearance and the plan as a planning activity. Then, as we already did in (Doublait & Lelouche, 1989), we shall call plan operationalization the process through which the plan is transformed into the actual retained sequence of elementary behaviors. These various processes are detailed in section 3.

3. MODELING ABSTRACTION AND ADAPTATION IN THE DECISION-MAKING PROCESS

In this section we study in more detail the decision-making process introduced in section 1. We assume that decisions are to be taken by the reference actor in order to adapt to recent changes in the environment at discrete points in time. This process is driven by implicit long-term[2] goals for each actor to reach in his/her existence (e.g. self-preservation goal, career goal, winning goal for an actor in the context of a game, etc.). These goals may be conflicting (e.g. winning a game) or not (e.g. self-preservation in most situations) with the goals pursued by other actors in the scene.

We divide the decision-making process into two consecutive phases: analytical and resolution. During the analytical phase, the reference actor is to generate and modify his/her mental attitudes based on his/her interpretation of recent observed data (background information or observable statuses and behavior of actors, including himself or herself) and on currently existing mental attitudes. During the resolution phase, the reference actor is to use these mental attitudes to explicit them into observable behaviors through goals and plans aimed at reaching these goals.

In 3.1. we examine the characteristics of various abstraction and refinement multi-level hierarchies that hold the information in the decision-making process. Then we illustrate in 3.2., through an example, how some transitions may occur between these hierarchies to infer additional informations and ensure continuity of the mental process.

3.1. Abstraction levels in the decision-making process

We now examine how we use abstraction as a key concept to organize the information in the two phases of the decision-making process shown in figure 1.

Mental attitudes or behaviors exist at different levels of abstraction. Different actors may reach different levels of abstraction depending on the complexity of the environment and on the actor's abstraction capabilities. Lenat et al (1979) consider that knowledge exists at different levels of detail and abstraction, and that the use of the appropriate level significantly enhances the expressiveness and efficiency of reasoning. According to Malec (1989), multi-level organization of knowledge processing modules allows the simplification of the resulting procedures.

Since the scope of our paper is very general, we give to «abstraction» a broad meaning, where «more abstract» may imply some or all of «more general», «less detailed», «less concrete», and consequently often «more subjective» and/or «more uncertain» The common denominator of all types of abstraction is always that whenever the abstraction level of some characteristic of property is raised, the set of cases for which this characteristic of property holds is raised also. Note that the opposite is not necessarily true; for example, changing «red» into «red or blue» does not raise the abstraction level (it merely changes a range of values at the same abstraction level), whereas changing «red» into «reddish» or «mostly red» (depending on the type of objects to which it applies and on the type of inference to be made) does increase the abstraction level.

3.1.1. *Variations in the analytical and resolution phases*

In our case, the nature of the analytical processes (from observable data to possibly high-level mental attitudes) is such that they are characterized by gradually increasing levels of abstraction; on the contrary, the resolution processes (from high level mental attitudes to observable actions) are characterized by decreasing levels of abstraction.

Both phases also vary in terms of information truthfulness. Indeed, the analytical processes infer information which is either true (knowledge) or uncertain (beliefs), whereas the resolution processes lead to information that is either true (knowledge) or intentionally false (bluff intentions).

Finally, both phases differ in terms of the time tags associated to the inferred information. If the decision-making process takes place between time points T0 and T1, all information inferred during the analytical phase characterizes the environment (background and actors) as it was at time points before T0. However, all information derived during the resolution phase characterizes actions to be taken by the reference actor after the decision-making process is completed, thus after T1.

Figure 3 illustrates the dimensions of abstraction, truthfulness, and time during one decison-making cycle, as described above. Two distinct hierarchies of abstration exist in each phase: descriptors and beliefs in the analytical phase, goals and plans in the resolution phase. We explain these terms in the next sub-section. High-level descriptors and beliefs are inferred from low-level ones by the means of (ascending) abstraction rules, whereas behaviors and low-level plans are inferred from high-level by the means of (descending) operationalization rules. No transition rule exists between goals of different abstraction levels, since goals are simply entry-points to the plans hierarchy (see 3.1.4.).

FIGURE 3.
Abstraction hierarchies in the decision-making process

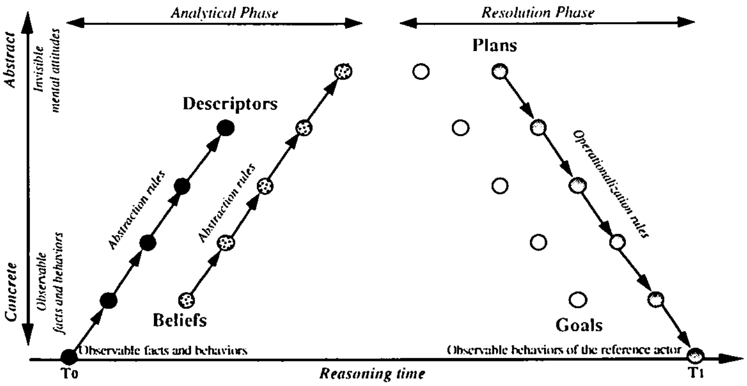

3.1.2. *Hierarchy of descriptors*

Descriptors refer to information that is certain, more or less abstract, and more or less subjective. They are derived only from observable facts and behaviors, and are refined in increasing levels of abstraction. For instance, «the current game turn is number 5» is a observable fact, whereas «the current game phase is the opening» is a more abstract descriptor derived from this fact. The latter is a descriptor since it is certain, although more abstract and also more subjective than the initial observable fact from which it is derived (observable facts and behaviors are objective by nature). In theory, there is no fixed upper limit in the abstraction of descriptors: that limit, if there is one, depends on the application.

In figure 4, we show which knowledge elements defined in section 2 can be modeled inside the descriptor hierarchy. We consider that these knowledge elements may apply to background information, as well as to the reference actor or to other actors. The latter distinction is justified by the fact that all abstraction hierarchies defined here are developed by one particular actor (the reference actor) who has by nature a complete and certain knowledge of himself or herself, as well as a partial (and possibly uncertain) knowledge of other actors.

FIGURE 4.
Knowledge elements of the descriptor hierarchy

DESCRIPTORS	*about the reference actor*	*about other actors*	*about the background*
Background statements	—	—	All
Statuses	All	Observable only	—
Behaviors	All	Observable only	—
Absolute roles	All	Observable only	—
Loaned roles	Ø	Ø	—
Borrowed roles	Ø	Ø	—

Since the reference actor knows everything about himself/herself by nature, all statuses, behaviors, and absolute roles of the reference actor are modeled inside the descriptor hierarchy at various levels of abstraction. However, only observable statuses, behaviors, and absolute roles of other actors can be modeled by descriptors, at various levels of abstraction. Typically, most other actors' behaviors are observable (e.g. «run»), some statuses are observable (e.g. physical statuses), and only few absolute roles are observable since they represent by definition mental, and thus hidden, relationships between actors (e.g. psychological, affective). By definition, loaned roles cannot be modeled by a descriptor since they lack certainty. A fortiori, borrowed roles cannot be modeled by a descriptor since they corerespond to actor intentions and obviously do not even belong to the analytical phase.

The assumption we made about the background (se 2.3.) implies that all background statements are observable and thus may be modeled by appropriate descriptors at various levels of abstraction.

3.1.3. *Hierarchy of beliefs*

Beliefs correspond to information that is uncertain (i.e. either true or false). They can also be expressed at various levels of abstraction. For instance, «there are few cookies in the cookie jar» is more abstract than «there are 3 cookies in the cookie jar». Both statements are beliefs (with possibly different associated levels of confidence), assuming for instance that the jar content is unknown but can be estimated by its weight. By definition, no belief is certain and observable, and thus the belief hierarchy in figure 3 starts at a higher level of abstraction than the descriptor hierarchy. In theory, like descriptors, there is no finite upper limit in the abstraction of beliefs.

Notice that in figure 3, the belief hierarchy is at he right of the descriptors hierarchy on the time scale. This is justified by the fact that beliefs are derived from descriptors and not the contrary, as we will explain it in section 3.2. For that reason, as also illustrated in figure 3, beliefs typically reach higher levels of abstraction than descriptors.

In figure 5, we indicate which knowledge elements can be modeled in the belief hierarchy. All information defining the reference actor is certain by nature, and thus cannot be modeled in the belief hierarchy. Similarly, all background statements are by definition certain and thus cannot be modeled in the beliefs hierarchy either.

Figure 5.
Knowledge elements of the belief hierarchy

BELIEFS	*about the reference actor*	*about other actors*	*about the background*
Background statements	—	—	Ø (our assumption)
Statuses	Ø	All (through loaned roles)	—
Behaviors	Ø	All (through loaned roles)	—
Absolute roles	Ø	All (through loaned roles)	—
Loaned roles	Ø	All (through recursive loaned roles)	—
Borrowed roles	Ø	All (through recursive loaned roles)	—

Since loaned roles define all beliefs of an actor regarding another actor, all of them are definitely modeled in the beliefs hierarchy at various levels of abstraction. In the simplest cases, the subject argument of a loaned role may be a status, a behavior, or an absolute role. However, since loaned roles are recursive (see 2.4.5.), the subject argument may also be a borrowed role or anther loaned role. We define the abstraction level of a loaned role as the abstraction level of its subject argument plus one. This definiton takes into account the intuitive fact that recusivity augments the abstraction (and the complexity) level: higher the recursion, higher the abstraction.

3.1.4. *Hierarchy of goals*

A goal is a state to be reached by the reference actor in order to adapt to (possibly recent) modifications of the environment. Goals are certain and true, although their actual attainment can only be assessed in the future (i. e. after the decision-making process ends). Goals can also be expressed at different levels of abstraction.For instance «being present where movies are made within a week» is more abstrat than «being present at Hollywood MGM studios in five days». As illustrated in these examples, time considerations are an inherent part of goals (either implicitly or explicitly) in the form of duration or specific dates by which the goal is to be attained.

Although goals may be defined at different abstraction levels, we assume no transitions between the various levels of the goal hierarchy, which explains the absence of arrows in figure 3. The reason is that, it makes no sense to decompose a goal into lower abstraction subgoals, since a goal refers to a state to be reached in the future; such subgoals would never reach the level 0 since goals cannot be observed. However, plans implementing these goals can ultimately be observed once they reach the level of directly executable actions (level 0 in the plan hierarchy). In fact, goals are simply entry-points to the hierarchy of plans (see 3.2.).

Goals can be set by the reference actor with the sole purpose of misleading other actors, thus being equivalent to the bluff intentions defined in section 2. For instance, «Have John believe that I am a movie star» is a goal which characterizes a false status of the myself (the reference actor) intended to be believed by some other actor(s), here John.

In figure 6, we indicate which knowledge elements can be modeled in the goal hierarchy.

FIGURE 6.
Knowledge elements of the goal hierarchy

GOALS	*about the reference actor*	*about other actors*	*about the background*
Background statements	—	—	Ø (our assumption)
Statuses	Yes, if attainable	Yes, if observable and reasonably attainable	—
Behaviors	Yes, if attainable	Yes, if observable and reasonably attainable	—
Absolute roles	Yes, if attainable	Yes, if observable and reasonably attainable	—
Loaned roles	(see figure 7)	(see figure 7)	—
Borrowed roles	(see figure 7)	(see figure 7)	—

Since goals represent states of the (reference actor's) universe to be attained, all knowledge elements that can be modeled by descriptors can also be modeled by goals, as long as the states thus described can reasonably be attained. As we saw in 3.1.2. and in figure 4, those include statuses, behaviors, and absolute roles, describing other actors as well as the reference actor. Examples of goals involving only the reference actor are: («I want to») «be a fireman» (a social status), «protect myself from the rain» (an abstract behavior), or «be Mary's husband» (an absolute relational and affective role). A goal involving another actor might be: «I want my son to become a lawyer» (a social status).

Moreover, since a state of the universe may include beliefs and bluff intentions, i.e. in our case loaned and borrowed roles, those can also be goals (at least partial). However, such roles cannot be shown in figure 6 as they were in figures 4 and 5, because goals may also involve the other actors' universes, and not only the reference actor's one. If we want to adequately model a goal involving loaned or borrowed roles, we must distinguish the owner of the role (i.e. in which universe it belongs) and whom this role concerns (i.e. at whom it is directed). This is done in figure 7, which clearly shows that I cannot have as a goal to believe something (roles which would be loaned by the reference actor) or to make myself believe something (borrowed roles which would be directed at the reference actor), but that I certainly can have as a goal to make somebody else believe something (roles loaned by others, through conviction), and most important, to act (or to convince an accomplice to act) in such a way as to induce some (possibly false) belief in someone else's mind.

FIGURE 7.
Loaned roles and borrowed roles as goals

Goal elements of the reference actor	*by the reference actor*		*by another actor*	
	about the reference actor	*about other actors*	*about the reference actor*	*about other actors*
Role loaned	No (meaningless)		Yes (through conviction)	
Role borrowed	No[3]	Yes, all	No[4]	Yes (through conviction)

The last case is the new element brought here: borrowed roles of the reference actor can only be part of the resolution phase, and (even more importantly) can only be goals, at various levels of abstraction. As for loaned roles, we may define the abstraction level of a borrowed role as the abstraction level of its subject argument plus one.

3.1.5. *Hierarchy of plans*

Plans represent more or less abstract actions which, together and in the proper sequence, can implement a particular goal. Plans can be recursively decomposed into sub-plans of lower abstraction, until these sub-plans cannot be refined (or de-abstracted) anymore. A sub-plan that reaches its refinement limit corresponds to an action which can be directly executed by an actor, and thus can be observed by the entire actor community if it is a public action. The process of refining a plan into lower level plans is called operationalization. A goal can be implemented by different plans. We shall say that a plan has abstraction level N if it implements a goal of identical abstraction level N.

In fact, a plan can be represented by a n-ary tree as illustrated by the partial example in figure 8. Each level of the tree represents a different level of abstraction, and is ordered to respect the plan sequence. Leaf nodes are atomic actions (at level 0) whereas intermediate nodes are sub-plans requiring further refinements.

Since a plan is initially an abstract action and ultimately a structured sequence of actions, it may not contain background statements, statuses or roles, but only behaviors. Moreover, since a plan implements a goal of the reference actor, it may only contain behaviors planned to be executed by the reference actor. In other words, the plan hierarchy contains only behaviors of (i.e. executable by) the reference actor, at various levels of abstraction.

FIGURE 8.
Tree representation of a high-level plan

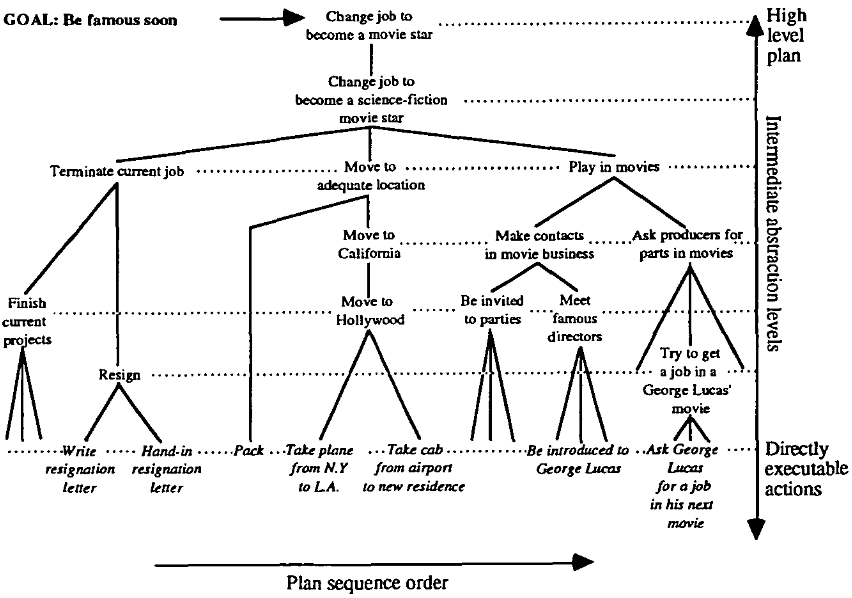

3.2. Transitions between abstraction hierarchies

In 3.1., we showed how descriptors, beliefs, goals, and plans can be defined at various levels of abstraction. We also showed the existence of intra-hierarchy transition rules, namely abstraction rules to derive higher-level descriptors and beliefs from lower-level ones, and operationalization rules to refine high-level plans into lower-level ones and ultimately into observable behaviors.

The four hierarchies described in 3.1. also interact with each other by the means of interhierarchy transition rules. These rules are categorized as follows:

- belief generation and adaptation rules to operate transitions from the descriptor hierarchy to the belief hierarchy,
- goal generation rules to operate transition for the descriptor and/or the belief hierarchies to the goal hierarchy,
- plan generation rules to operate transition from the goal hierarchy to the plan hierarchy.

These rules are necessary to maintain the continuity of the decision-making process. They allow for the adaptation and evolution of mental attitudes in the model of the reference actor. The next paragraphs detail each of these rule categories, using complete examples.

3.2.1. *Belief generation and adaptation rules*

Belief generation rules induce beliefs from descriptors and possibly observable facts. More than one descriptor may be needed to induce a single belief, and these descriptors, may reside at different levels of abstraction, either lower or higher (but more likely lower) than the induced belief.

Figure 9. Illustrates four consecutive reasoning cycles of the reference actor (referred to as «I» in the example), yielding four steps in the overall analytical reasoning process. Underlined statements correspond to information observed or derived in the current step.

In step 1, the reference actor observes a behavior executed by actor John («John is attending a class»). This primitive behavior is then abstracted by and abstraction rule to infer the status descriptor «John is a student».

In step 2, the reference actor observes a behavior executed by actor Bill («Bill is teaching a class»). From this behavior, it induces a new abstract status descriptor («bill is a professor»). Descriptors defined previously are still present in the reference actor's knowledge base.

In step 3, the reference actor observes a new behavior of John: «John is giving a tap on bill's shoulder». This primitive behavior is then abstracted by an abstraction rule to the higher-level absolute role descriptor «John has physical contacts with bill». From this level of abstraction, another abstraction rule induces an even higher-level psychological absolute role descriptor («John is familiar with Bill»). Finally, based on the overall abstract knowledge that "John is a student», «Bill is a professor», and «John is familiar with Bill», the reference actor induces the belief (a loaned role) that «Bill is John's professor». This loaned role has three arguments: the reference actor (owner of the loaned role), Bill (the actor to whom the role is loaned), and the social absolute role that defines bill as John's professor (the subject of the loaned role). Since the abstraction level of a loaned role is defined as the abstraction level of its subject argument plus one, «I believe that bill is John's professor» should be one level higher than the potential descriptor «John is Bill's professor» (which is at the same level than «Bill is a professor» and «John is a student»). Note that the expression of a loaned role is not unique, since the reference actor equivalently believes that «John is Bill's student».

FIGURE 9.
Belief generation and adaptation rules

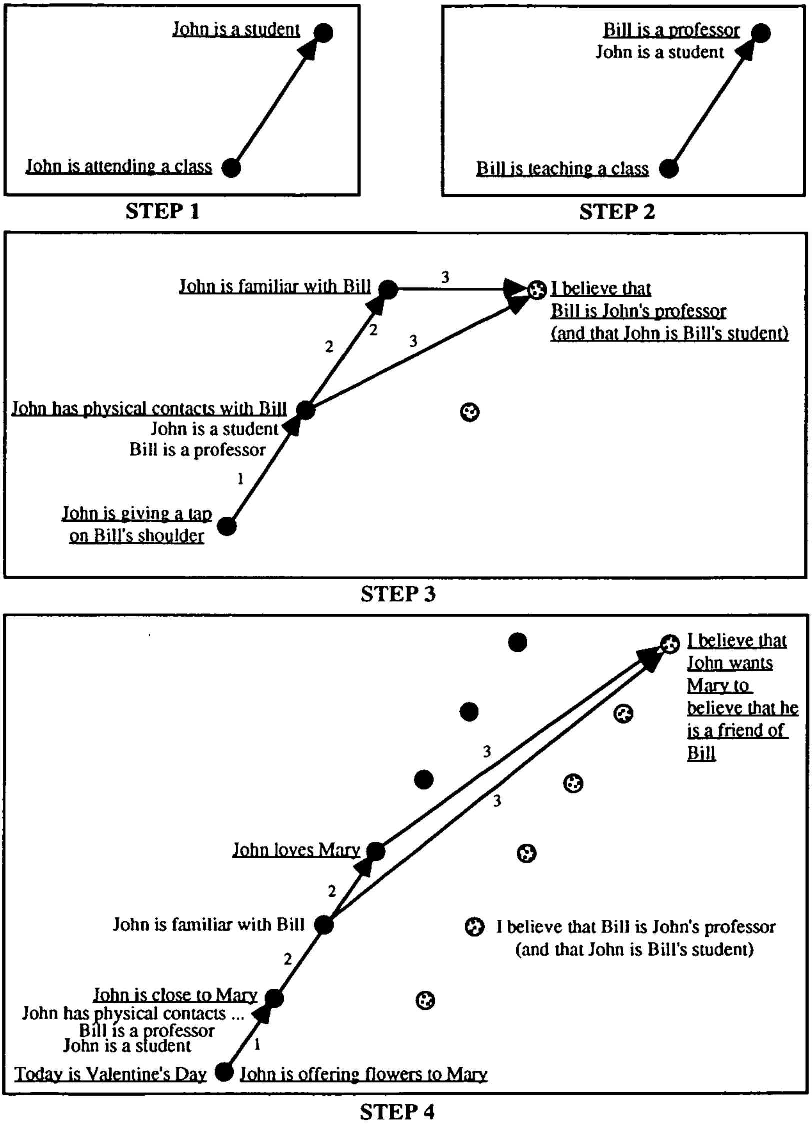

In step 4, the reference actor observes another behavior of John («John is offering flowers to Mary»). It also knows the fact that «Today is Valentine's Day». Based on these two observable knowledge elements and possibly other similar ones, the reference actor may abstract the high-level absolute role descriptor «John is close to Mary». Then, this descriptor can be further abstracted into the higher-level descriptor «John loves Mary», which is at a higher level than «John is familiar with Bill» (indeed, «John loves Mary» should be at the same level as the potential descriptor «John is a friend of Bill» (indeed, «John loves Mary» should be at the same level as the potential descriptor «John is a friend of Bill», which itself should be higher than «John is familiar with Bill»). Finally, based on this last descriptor and the previous abstrat evidence that «John is familiar with Bill», the reference actor may induce the belief that «John wants Mary to believe that he is a friend of Bill», assuming that John wants to impress Mary in order to win her heart. This belief is a recursive loaned role with three arguments: the reference actor (owner of the loaned role), John (the actor to whom the role is loaned), and the borrowed role «John wants Mary to believe that he is a friend of Bill» (subject of the loaned role). This borrowed role is in turn made of three arguments: John (owner of the borrowed role), Mary (actor who is the target of the borrowed role), and the absolute role expressing the friendship betwen John and bill, Based on the definition of the level of abstraction of a loaned role, and on the recursively computed level of its subject argument, this derived loaned role must be four levels higher than the descriptor «John is familiar with Bill». Indeed, if «John is familiar with Bill» is at level N, the potential descriptor «John is a friend of Bill» is at level N+1, the potential belief «Mary believes that John is a friend of Bill» is at level N+2, the potential goal (borrowed role) «John wants Mary to believe that he is a friend of John» is at level N+3, and finally, the belief «I believe that John wants Mary to believe that he is a friend of John» is at level N+4.

As we illustrate in this example, potentially conflicting beliefs may co-exist in the same hierarchy. Indeed, four situations are possible, among which only one is true. 1) John is Bill's student and has no intention to impress (and lie) to Mary, or 2) John is Bill's student and wants to impress Mary, or 3) John is not Bill's student but wants to impress Mary, or finally 4) John is not Bill's student and does not want to impress Mary, and John's behavior is driven by knowledge which is still unknown to the reference actor. Situations 1) and 3) are conflicting, whereas situation 2) is positively supportive, and situation 4) is negatively supportive. Conflicting beliefs are useful since they provide adaptation opportunities that may lead to changes in the rule base in order to avoid further conflicts of the same nature.

Belief abstraction rules as well as belief generation and adaptation rules derive conclusions that are more powerful (to justify their existence), but also less certain, than their premises. However, the multiplicity of independent rules that lead to the same conclusion, and to a certain extent the number of premises per rule, may increase the certainty of this conclusion.

3.2.2. *Goal generation rules*

Taking the same love story between John and Mary as an example, we now focus on the transition between the analytical phase and the resolution phase. This transition is achieved by goal generation rules. Every such rule uses descriptors and beliefs at various abstraction levels to induce a goal at a particular abstraction level.

More than one goal may be induced in the same reasoning cycle. In such a case, a selection must be made to determine a single active goal which will be further implemented by a plan. As usual in rule-based systems, various selection strategies can be used, such as associating a weight to each goal, which allows for adaptation of goal generation rules to select the best possible goal.

In figure 10, we asume that the reference actor suddenly falls in love with Mary and thus directly infers the high-level descriptor «I love Mary» (at the same level as «John loves Mary»). Since this is an absolute role that belongs to the reference actor, it can be derived from some hidden (non-observable) personal behaviors (or feelings), and the abstraction rules that induced this descriptor may be omitted. Simultaneously (i.e. at the same discrete point in time), the reference actor is observing Mary's behavior («Mary is ignoring John», where «ignoring» is taken in the concrete sense here, i.e. she pretends that she does not see him). From this behavior, and knowing already that today is Valentine's Day and that John just offered flowers to Mary, the reference actor can induce the high-level descriptor «Mary does not like John» or even «Mary hates John» (at the same level as «I love Mary» and «John loves Mary»). The reference actor is willing to go further in his reasoning process and induces a very high-level belief from the descripors: «Mary hates John» and «John is familiar with Bill». The twice recursive loaned role believed by the reference actor is: «Mary believes that John wants her to believe that he is a friend of Bill». The owner of this loaned role is the reference actor. The actor to whom this role is loaned is Mary. The subject is another loaned role, owned by Mary, loaned to John, and whose subject is a borrowed role (owned by John, towards mary, and whose subject is the absolute role defining the friendship between John and Bill). The belief inferred at step 5 has a higher abstraction level than the belief inferred at step 4 since the degree of recursively nested roles is greater (see definition in 3.1.3.).

FIGURE 10.
Goal generation rules

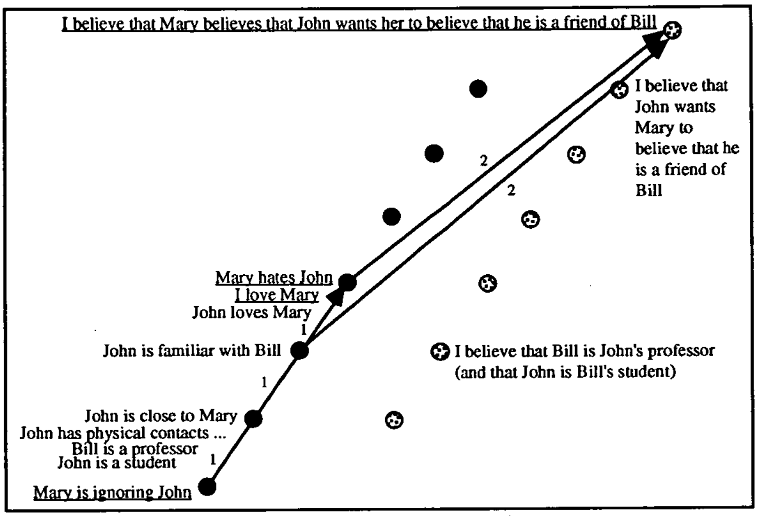

STEP 5 (last analytical step)

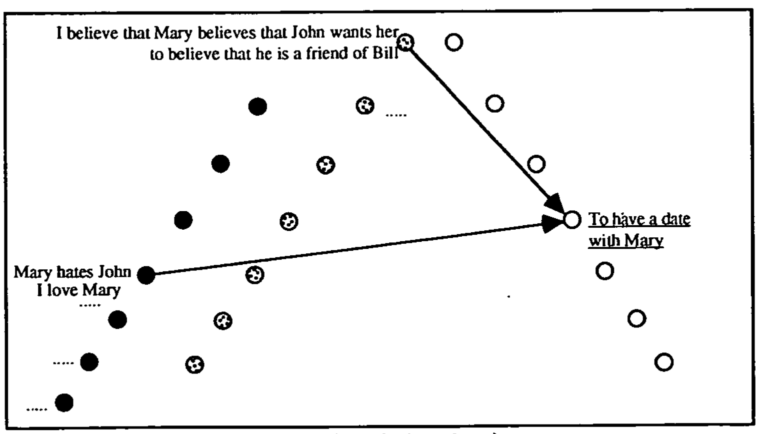

STEP 6 (Resolution phase)

As explained at the end of 3.2.1., this last belief is likely to have a relatively low level of confidence considering the absence of further evidence, and the fact that it has a much higher level of abstraction than its premises. Highly confident beliefs are typically induced from premises at similar abstraction levels, preferably higher.

Assuming that step 5 completes the analytical phase, the reference actor enters the resolution phase. He must then determine a goal that will allow him to adapt to the newly analyzed environment. Since the reference actor knows that he loves Mary, and assumes that Mary believes that John is a jerk who simply wants to impress her, he is gaining confidence and decides to have a date with Mary to adapt to his new feeling and to the situation change. This is done in step 6, where the second diagram of figure 10 illustrates the particular generation rule that induces the «dating» goal.

Goals can be induced from descriptors only, from beliefs only, or from a combination of descriptors and beliefs as illustrated in this example. In all cases, premises can exist at any abstraction level. It is important to note that there is no possible mapping between abstraction levels in the analytical phase (premises of goal generation rules) and abstraction levels in the resolution phase (conclusions of goal generation rules). Indeed, these are two totally disjoint universes that express knowledge of a different nature.

3.2.3. *Plan generation rules*

Once a particular goal has been selected, a plan must be chosen to implement it. Planning is a complex activity, and we will not address this topic in detail in this paper (for further information on planning, the reader may refer to [Schank & Abelson, 1977]). To simplify our discussion here, we will assume the existence of a rule-based planning mechanism in which each plan generation rule takes on a single goal premise and induces one or more high-level plans of the same abstration level (as proposed in section 3.1.5., a plan of abstraction level N is defined as a plan implementing a goal of identical abstraction level N).

Figure 11 illustrates the process by which the goal selected by the reference actor («to have a date with Mary») is assigned by the planner in step 7, a plan of equal abstraction level: «to do something unique for Mary».

FIGURE 11.
Plan generation rules

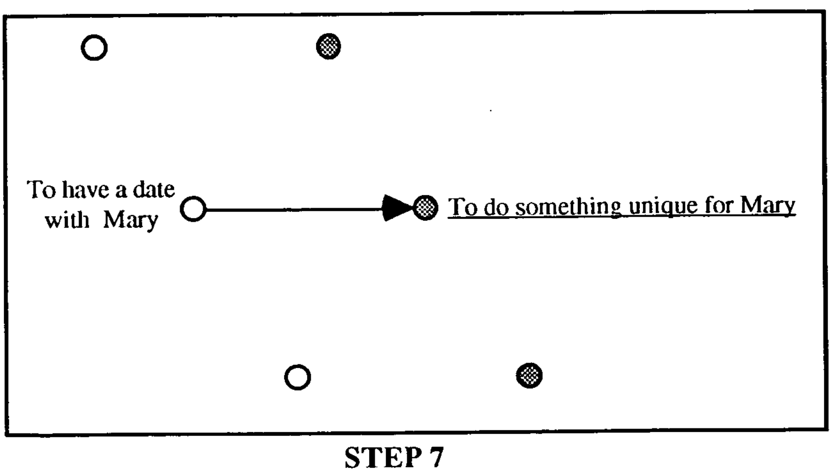

The final step of the resolution phase would consist in using operationalization rules to refine the high-level plan and gradually decompose it into a sequence of atomic plans, i.e. actions directly executable by the reference actor. For instance, a sequence of subplans to be executed by the reference actor might be: 1) «to buy a ladder», 2) «to learn a melody», 3) «to use the ladder to climb to Mary's window», 4) «to sing the melody to Mary», 5) «to observe Mary's reaction», and 6) «to ask mary out if she reacts positively». The operationalization process is discussed in more detail in 3.1.5.

If this plan proves to be unsuccessful (after subplans 5 or 6), Mary's reactions may prompt changes in the reference actor's mental attitudes (beliefs, goals, or plans) at a next decision-making cycle, so that another goal and/or plan may be selected. As long as it proves to bee successful (here after subplans 1 through 4, and possibly 5 and 6), the reference actor may elect to continue the active plan and directly execute the next planned subplan, thus by-passing all goal and plan generation rules.

4. HUMAN-ASSISTED MECHANISM FOR INTERACTIVE LEARNING AND ADAPTATION: A CASE STUDY

The reasoning cycle described in section 3 allows for adaptation through the modification of transition rules. However, we still have to explain how adaptation is performed in the cycle. Furthermore, learning new rules is another cognitive process which is also required to refine and correct the reference actor's model of the universe. The present section is aimed at addressing these questions.

In order to experiment our research on bluff, we adopted a strategy game as the problem domain. Indeed, complex game may be ideal laboratories for the study of some cognitive processes such as abstraction, adaptation, and learning. In addition, the game we selected offers a great deal of similarities with the multi-agent model just defined, even including bluff as one of its major characteristics. We briefly introduce this game in section 4.1. In 4.2., we study common approaches to solve theory refinement problems. Finally, 4.3. presents a proposal for a human-assisted learning mechanism in the context of Sigma-File.

4.1. The Sigma-File game

Our domain game is Sigma-File, a four-player strategy game. The full game description is in (Parlett, 1977). The playing board consists of several cities connected to one another by various routes. There are nine game tokens: eight independent international spies which can be moved around by all players, and one secret file that is lured by four intelligence services (KGB, CIS, CIA, and IS), each one being located in one of the board cities (Moscow, Peking, Washington, and London respectively) and headed by one of the four players (see figure 12). The common (conflicting) goal of every player is to bring the secret file to his own capital.

FIGURE 12.
The Sigma-File game board.

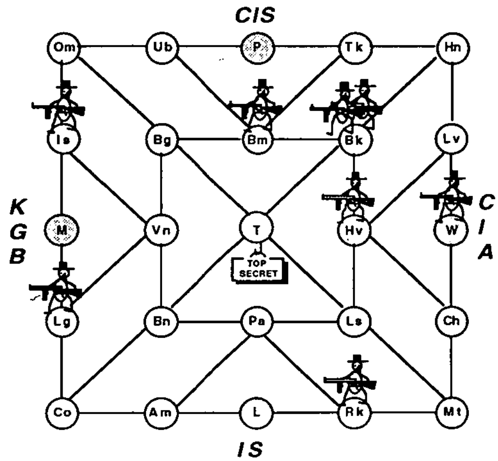

At the start of the game, the secret file is located at the city at the center of the board (Tangier), and every player has at his disposal an initial cash balance (10,000 dollars) that he can use during the game to pay one or several spies and thus attempt to control them.

When his turn comes, a player may choose between three possible moves: to transfer a spy from its current city to and adjacent one (the spy may also carry the file along if it is located in the same city), to pay a spy a certain amount of money which is then subtracted from the player's current balance (the spy's identity and the amount are both kept secret from his opponents), or to have a spy killed by another spy on which he has at least a 1000-dolar control (that fixed amount is subtracted from the current control amount of the player on the murderer). Finally, since no player exactly knows who controls which spy, transfers and murders may be challenged by any opponent to prevent their execution. Theoretically, the player who has the highest control amount on the implied spy wins the challenge, but the challenger may choose to either reveal himself (if he really wants to stop the current player's move) or to bluff his opponents (if he deliberately looses the challenge). A same move may be challenged by several opponents in turn. In any case, each spy remains loyal to whoever admits to be paying it most at any given moment.

4.2. Various learning paradigms to handle theory refinement

The process of gradually augmenting and correcting an approximate or incomplete knowledge base is often referred to as theory refinement (not to be confused with the refinement of a plan introduced in section 3, which is the process of reducing a plan defined at a high level of abstraction into a series of less abstract subplans) as explained in (Ourston & Mooney, 1990).

A theory can be defined in terms of positive and negative examples. It is correct and complete when no negative examples can be proven and all positive examples are provable in this theory. Based on this statement, Ourston and Mooney (1990) define two forms of incorrectness in a theory: over-generality and over-specificity. Incorrect theories can have both over-general and over-specific aspects.

Over-generality occurs when some negative examples are provable. Two types of errors can lead to over-generality: an incorrect rule is present in the theory, and/or an existing rule lacks some constraint in its premise. The solution consists here in removing rules, restricting constraints, or both.

Over-specificity occurs when some positive examples are not provable. Two types of errors can lead to over-specificity: a rule which is necessary to the proof of the certain examples is missing from the theory, and/or an existing rule has additional constraints in its premise. The solution consists here in adding rules, loosening constraints, or both.

Typical approaches to address the issue of theory refinement use explanation-based learning or empirical learning methods. Explanation-based learning (EBL) is a form of inductive learning, whereas empirical learning covers a wide range of methods, including the learning by instruction paradigm (see [Michalski & al., 1983]). EBL is efficient at detecting errors and focusing corrections on the failing portion of the theory, but empirical methods are usually efficient in modifying and adding new rules. Furthermore, empirical methods often lead to faster results than purely inductive ones. Theory refinement is not a trivial issue, and most approaches handle either over-generality only (Flann & Dietterich, 1989), or over-specificity only (Wilkins, 1988; Ali, 1989). On the other hand, handling multiple faults in a theory is more complex than handling single faults, and some approaches are restricted to the latter. Ourston and Mooney (1990) developed an approach that combines both EBL and empirical methods to address both over-specificity and over-generality, and also deal with multiple faults.

4.3. Human-assisted learning in Sigma-File

4.3.1. *Learning scenario*

Most true multi-agent environments, such as Sigma-File, have models which are significantly more complex than most micro-worlds addressed by EBL techniques. Indeed, complex (sometimes recursive) beliefs, bluff intentions, and conflicting goals, almost impose a human-assisted type of learning to ensure proper monitoring of acquired and modified knowledge. In addition, empirical methods are usually more efficient.

Therefore, we adopted a learning by instruction paradigm to model an automatic player playing Sigma-File against three human players (see [Doublait & al., 1987]). This automatic player (called AP for the sake of clarity) is the reference actor in whom[5] we are interested and whom we are trying to model. Indeed, AP has the capability of observing and analysing his opponents' moves. Since he has almost no status and role information regarding his opponents, he must build beliefs about them. He has also full access to background information (game turn, positions of spies on the board, etc.) like any other actor, as well as complete knowledge of his own statuses (balance, bids on spies). Finally, AP is capable of bluffing his opponents by pretending false statuses and absolute roles.

We intentionally decided to load AP's knowledge base with only limited initial skills, thus making our learning problem fall into the theory refinement category. AP knows about the rules of Sigma-File, and basic strategies defined in textbooks. Our goal is to have AP gradually improve his moves and thus play better. We are not as interested in the final skill level AP can reach as we are in the progress he can accomplish from his initial level.

In our approach, AP plays against his human opponents under the constant supervision of a human tutor whose responsibility is to observe AP's derived knowledge and to provide feedback on the correctness and relevance of his decisions in the current situation. Once an error is detected in AP's reasoning cycle, the tutor's objective is to enter a learning and adaptation dialogue with AP to isolate and correct the error. We intend to handle both over-specificity and over-generality as well as multiple faults.

4.3.2. *Identification of intervention break-points*

The tricky point consists in determining the break-points in AP's reasoning cycle when the tutor can intervene. Ideally, the tutor should be allowed to interrupt the process each time a transition rule is executed. However, it is impractical to do so in a real-time learning environment such as ours, because we want to keep the game running. Instead, we defined several fixed intervention break-points in the reasoning cycle. We believe, although we did not prove it yet, that the results obtained should be overall as good as with a step-by-step execution mode. The price to pay for having fewer break-points is that, if an error is detected, the source of the error may be located far behind in the reasoning chain, thus forcing expensive backtracking activities to trace the reasoning cycle until at most the previous break-point. Consequently, we still try to keep a short distance between two consecutive break-points.We decided to have the tutor give feedback to AP as indicated in figure 13.

4.3.3. *Nature of the tutor's intervention*

At each defined break-point, the tutor is allowed to observe the conclusions derived by AP up to this point, and take one of the following actions:

1. Invalidate a conclusion (over-generality situation; a negative example has been proved), that is: a) an incorrect transition rule that derived this conclusion is present, and/or b) an existing transition rule that derived this conclusion is missing a satisfied constraint in its premise.
2. Indicate that a conclusion is missing (over-specificity situation, not all positive examples are proved), that is: a) a transition rule is missing to derive the desired conclusion, and/or b) an existing rule has an unnecessary constraint which is not satisfied so the desired conclusion could not be reached.

At break-points where conclusions have an associated weight, such as beliefs (level of confidence), goals (priority of the goal relative to other candidate goals), and plans (priority of the plan relative to other candidate plans), the tutor should also be able to:

3. Invalidate the weight of a conclusion (either too low or too high), that is: a) an incorrect weight is asociated to one or more premises, and/or b) an incorrect weight is asociated to the transition rule itself.

When incorrect conclusions are derived (situation1) using more than one rule[6], the last executed rule may not be responsible for the error, but rather a previous rule is the inferencing chain may be the cause. In this case, the tutor needs to backtrack AP's search, rule by rule, from the last executed one, until he reaches the rule which is responsible. Likewise in situation 2, a correct rule might not have been fired because one of its premises (which is satisfied) has not been derived by a previous missing or incorrect rule. A similar situation may arise in situation 3, where the weight of a conclusion is incorrect due to giving an inadequate weight to an earlier rule in the inferencing chain.

FIGURE 13.
Tutor's intervention break-points in the reasoning cycle

Logical break-point	*Necessity for a tutor's intervention*
After a descriptor abstraction rule is fired	No. There is no need for the tutor to intervene after each descriptor abstraction rule since by definition these rules preserve the truthfulness of the information. It is a relatively low-risk transition.
After a belief abstraction rule is fired	No. For the same reason as above, abstraction rules preserve the level of certainty (here we should say the level of uncertainty) and thus these rules are less critical than inter-hierarchy rules.
After a belief generation or modification rule is fired	Yes. The risk is higher here, due to the uncertainty of the conclusion of these rules. Consequently, a closer monitoring of the reasoning is required for belief generation and adaptation rules. The tutor should verify that all derived beliefs are acceptable in the current context, that no obvious belief is missing, and that the associated level of confidence of the belief is also adequate.
At the end of the analytical phase	Yes. Obviously, reconciling the analysis (descriptors and beliefs) of the situation is required to validate the entire phase before starting the resolution phase.
After a goal generation rule is fired	No. Since the tutor intervenes after all goal generation rules are fired, there is no need to duplicate the effort.
After all goal generation rules are fired	Yes. The tutor needs to validate all derived goals to determine whether these goals are all correct or some are missing. Also, since there is a selection process involved, he should verify the goal classification to ensure that the goal priorities are correct.
After a plan generation rule is fired	No. Since the tutor intervenes after all plan generation rules are fired, there is no need to duplicate the effort.
After all plan generation rules are fired	Yes. The tutor needs to validate all derived plans to determine whether these plans are all correct or some are missing. Also, since there is a selection process involved, he should verify the plan classification to ensure that the plan priorities are correct.
After an operationalization rule is fired	No. For reasons similar to abstraction rules in the analytical phase, operationalization rules are low-risk transitions.
At the end of the resolution cycle	Yes. This final intervention will ensure that the plan decomposition is correct, and that the next moves to be executed are relevant to the current situation and consistent with the active goal and high-level plan.

4.3.4. *Type of dialog between the tutor and AP*

The quality of the interaction between AP and his tutor is critical in the success of the theory refinement process. A major problem consists in providing the tutor with an interface that is expressive enough to let him provide his feedback to AP in yet a simple and concise manner. Typicalliy, learning-by-instruction systems have a restricted natural language interface like in TEIRESIAS (Davis & Lenat, 1982), or a WIMP7-type interface, like in ASK (Gruber & Cohen, 1989). In our earlier work, we selected the latter approach for its simplicity, conciseness and user-friendliness aspects.

5.CONCLUSION

A large number of knowledge elements need to be modeled in a true multi-agent environment, including spatio-temporal facts as well as actors' characteristics, beliefs, goals, and plans. Modeling this information requires formalisms able to represent mental attitudes that may be true, uncertain, or even intentionally false in the case of bluff. Bluff is a very important concept that is inherent to many multi-agent environments, but surprisingly, that has rarely been modeled explicitly in AI (a notable exception is Waterman's [1970] POKER system).

We proposed a logical formalization of this type of information that allows the modeling of complex nested mental attitudes. Moreover, we extended our approach by demonstrating how the modelled knowledge elements could be represented at various levels of abstraction and how an actor could navigate through these abstraction levels to adapt himself or herself to recent changes in the environment. More importantly, we indicated how these complex models can be refined through human-assisted learning.

These ideas have already been partially implemented in the TACOS system (Tactics Acquisition an Operationalization System). We are currently working on the re-design and implementation of the concepts presented in this article using an object-oriented approach (cf. [Rumbaugh & al., 1991]). We believe that designing a multi-agent system using an object-oriented approach greatly eases its implementation and maintenance. Indeed, the object-oriented paradigm offers powerful representation and reasoning mechanisms (e.g. inheritance, aggregation, encapsulation) for the modeling of actor characteristics and mental attitudes. We also wish to demonstrate how abstraction levels can be mapped onto all three key components of an object-oriented analysis model as defined by Rumbaugh and his colleagues, namely the object model, the dynamic model, and the functional model.

From a more theoretical point of view, we are simultaneously working on fully integrating beliefs and bluff intentions concerning the background in our general multi-agent model.

References

ALI, K. (1989). «Augmentinfg domain theory for explanation-based generalization», Proceedings of the Sixth National Conference on Machine Learning, Ithaca, NY, June 1989, 40-42.

BARNDEN, J. (1989). «Towards a paradigm shift in belief representation methodology», Journal of Experimental and Theoretical Artificial Intelligence, 1, 133-161.

DAVIS, R.; LENAT, D. (1982). Knowledge-based systems in artificial intelligence. McGraw-hill, New York, NY.

DOUBLAIT, S.; LELOUCHE, R. (1989). «Using abstraction levels to integrate the tactics acquisition process in a problem-solving task», Proceedings of the Second International Symposium on Artificial Intelligence, Centro de Inteligencia Artificial, Instituto Tecnológico y de Estudios Superiores de Monterrey, Monterrey, NL, Mexico.

DOUBLAIT, S.; LELOUCHE, R.; JANTA-POLCZINSKI, M. (1897). «Vers un système d'apprentissage automatique pour le jeu Sigma-file (Towards an automatic learning system for the Sigma-file game)», Informatique Cognitive des organisations, Proceedings of the ICO'89 Conference, 259-283, L'Interdisciplinaire Informatique, Lyon, France, June 1989.

FETZER, J. (1985). «On defining "knowledge"», Al Magazine, 6, 19 (Spring).

FLANN, N.; DIETTERICH, T. (1989). «A study of explanation-based methods for inductive learning», Machine Learning, 4(2), 187-226.

GRUBER, T.; COHEN, P. (1989). «The design of an automated assistant for acquiring strategic knowledge». SIGART Newsletter: special issue on knowledge acquisiton, N.° 108, april 1989, 147-150.

LELOUCHE, R., DOUBLAIT, S (1992). «Qualitative reasoning with bluff and beliefs in a multi-actor environment», International Journal of Man-Machine Studies, 36, Academic Press Limited, 149-165.

LELOUCHE, R. (1986). «Quelques aperçus sur la modélisation de la description statique d'une scène» (Some issues on modeling a scene static description). CIIAM 86, Proceedings of the second International Conference on Artificial Intelligence, Marseille, France, December 1-5, 121-139.

LENAT, D., GUHA, R. (1990). Building Large Knowledge-Based Systems, Addison-Wesley, Reading MA, 372 pp.

LENAT, D.; HAYES-ROTH, F.; KLAHR, P. (1979). Cognitive Economy, Working paper HPP-79-15, Stanford Heuristic Programming Project, Stanford University (Stanford, Calif.), June 1979, 46 pp.

MALEC, J. (1989). «Knowledge elicitation during dynamic scene description». SIGART Newsletter: Special issue on knowledge acquisiton, N.° 108, april 1989, 162-163.

MICHALSKI, R.; CARBONELL, J.; MITCHELL, T. (1983). Machine Learning: An Artificial Intelligence Approach, Volume 1, Tioga Publ. Corp., Palo Alto (California), 572 pp.
OURSTON, D.; MOONEY, R. (1990). «Changing the Rules: A Comprehensive Approach to Theory Refinement», Proceedings of the Eighth National Conference on Artificial Intelligence, AAAI-90, Boston, MA, July 29-August 3, Volume 2, Machine Learning section, 815-820.
PARLETT, D. (1977). Popular Indoor-Games, Coles Publishing.
RUMBAUGH, J.; BLAHA, M.; PREMERLANI, W.; FREDERICK, E.; LORENSEN, W. (1991). Object-oriented modeling and design, Prentice Hall, Englewood Cliffs, NJ, 500 pp.
SCHANK R.; ABELSON, R. (1977). Scripts, Plans, Goals and Understanding, Lawrence Erlbaum Assoc., Hillsdale, NJ.
WATERMAN, D. (1970). «Generalization learning techniques for automating the learning of heurisstics», Artificial Intelligence, 1, 121-170.
WILKINS, D. (1988). «Knowledge base refinement using apprenticeship learning techniques», Proceedings of the Sixth National Conference on Artificial Intelligence, ASAAI-88, St.-Paul, MN, August 1988, 646-651.
WILKS, Y. (1985). Relevance, Points of View and Speech Acts: An Artificial Intelligence View, Research Paper MCCS-85-25, 24 pp., Computing Research Laboratory, New Mexico State Unviersity, New Mexico.

Notes

[1] Acting upon the environment may also requiere a long sequence of elementary behaviors, but the effect of such behaviors is more easily predictable than if the goal is to make someone believe something false.

[2] The term of a goal is always relatively long with respect to the scene sequence that we are trying to describe.

[3] Psychological studies depict examples where the reference actor "wants to deceive himself or herself" for some ulterior motives which explain an apparently inconsistent or illogical behavior, but such "goals" are not the result of a conscious reasoning activity aimed at attaining a given state of the universe, and therefore are not goals in the sense of this paper.

[4] Similar considerations hold for a situation where the reference actor "wants someone else to deceive him or her".

[5] Since AP is an actor, we use the masculine gender to refer to him, and the neutral one for spies and other tokens.

[6] In the following rule sequence: $A \rightarrow B$, $B \rightarrow C$, and $C \rightarrow D$, C and D are derived respectively using two and three rules, whereas B is derived using only one rule.

[7] Windows, Icons, Menus, Pictures.