A Framework for Multi-Agent Belief Revision

Part II A Layered Model and Shared Knowledge Structure

Wei Liu[*] and **Mary-Anne Williams**[†]

Abstract

In this paper, a framework of a belief revision agent in a multi-agent environment is presented. In the agent knowledge kernel, by separating domain knowledge from social knowledge, various level of belief revision capabilities has been modularized into multiple interactive layers. A vertical 4-layer 2-pass architecture is adopted for the BR engine. Therefore, the framework can model various levels of belief revision, including Single Belief Revision, BR using information from Multiple Sources and Multi-Agent Belief Revision. In order to achieve multi-agent belief revision in a heterogenous society, a Shared Knowledge Structure(SKS) is proposed, which allows the sharing of knowledge as well as protecting private knowledge. Using graded knowledge in the SKS, the process of multi-agent belief revision is implemented using knowledge migration, which is the procedure that reshapes the knowledge structure triggered by the new information.

Introduction

Belief Revision (BR) techniques for single agents have been implemented successfully(SATEN) and applied to significant areas of application, eg. requirements engineering and marketing research(Williams 1998).

This paper provides a framework for the necessary infrastructure to extend existing single BR implementations to a multi-agent environment. In particular, it supports *BR in a Single agent environment*, *BR using information from Multiple Source* and *Multi-Agent Belief Revision*, which can involve collaborative team-based BR.

Multi-agent belief revision builds on the research effort towards Multi-Agent Systems (MAS). Part I(Liu & Williams 1999) of this paper provides a survey of multi-agent belief revision, and it provides a hierarchy of different types of BR available in a MAS, shown in Figure 1. An agent's BR ability is thus enhanced from the basic capability of handling its own observations to the ability of participating in the required joint effort

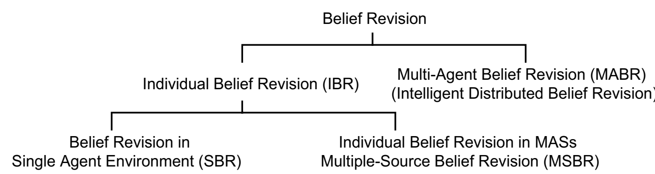

Figure 1: Belief Revision Hierarchy

for achieving society/group wide BR. In this paper, by separating *domain knowledge* from *social knowledge*, we show that the various levels of BR capabilities can be modularized into multiple interactive layers. A layered agent architecture is presented in next section, which serves as a flexible framework for various level of BR in MASs.

To tackle the social heterogeneity that might exist in knowledge systems, part I also demonstrates that a sophisticated knowledge structure is required to meet the following three requirements for a social agent[1]: (i) to share knowledge in order to cooperate with others or participate in teamwork, (ii) to keep some knowledge *truly* private in order to be competitive, (iii) to communicate and interoperate with other agents built on different structures adopted by previous researchers. In this paper, an ontological model of the agent knowledge structure, ie, *Shared Knowledge Structure*(SKS) is developed for such purpose. The proposed SKS is compared with other related structure in the same section.

As global consistency is not a requisite in SKS, one of our major observation is the *inconsistency principle.* Base on this principle and *knowledge grade* defined by SKS, the process of multi-agent belief revision turns into *knowledge migration.* This is investigated in a separate section. The paper concludes with future work in the last section.

A Layered Model of a BR Agent

BR is one of the basic skills of *deliberative agents*(Müller 1998). Excluding the domain capabilities, defined by

[*]Email: mgwl@infosystems.newcastle.edu.au

[†]Email: maryanne@infosystems.newcastle.edu.au

[1]A social agent is an agent that interacts with other agents in a MAS environment.

its general functions, a *BR agent* is an agent who can rationally change its knowledge base(s) in the face of new information.

Figure 2 illustrates a BR agent in a multi-agent environment, which consists of a perception subsystem, an action subsystem and a knowledge kernel. The knowledge kernel, which is the focus of this paper, contains a *social knowledge base*(SKB), a *domain knowledge base*(DKB), and a 4-layered BR engine.

The separation of cooperation knowledge (in our case, social knowledge) from domain knowledge is a theme that can be found in the early DAI work(Davis & Smith 1983). It leads to computational efficiencies and allows the agent's functions to be modularized into interactive layers. The contents of the SKB and the DKB is shown in the subdiagram of Figure 2.

The DKB provides the domain knowledge that agents use to reason, and it is the DKB that the single BR processes act on. Hence, for an agent to perform BR, there must, at least, exist a DKB upon which to carry out single BR if there is no SKB.

Compared to the stand-alone single agent, the SKB is one of the distinctive features possessed by agents who socially interact with others. By building operators to update such social information, the following properties would be achieved, (i) the *agent society*[2] would be open to any agent who joins in or leaves at any time, (ii) the agent would be able to handle polluted information, as is the case in Dragoni's BR using information from multiple sources(Dragoni & Guorgini 1999), and (iii) the agent would be able to work out the *candidate set of shared knowledge*, i.e. the *mutually accessible knowledge*.

The 4-layer BR engine based on a vertically layered two-pass flow architecture is the core of our BR agent. Vertical layering outperforms horizontal layering in terms of the feasibility of achieving coherent agent behavior(Wooldridge 1999). Two-pass control is more appropriate for our purpose because it allows feedback from higher levels to invoke the lower level functions, which in turn allows modeling heterogeneous agents with different levels of BR capability.

Inside the Layered BR Engine

Each layer has intra-layer functions and interfaces(consist of inter-layer functions) with adjacent layers. Intra-layer functions are responsible for processing the information contained in its layer. The inter-layer functions of the interfaces control the information flow from one layer to another and pass inputs and outputs to adjacent layers.

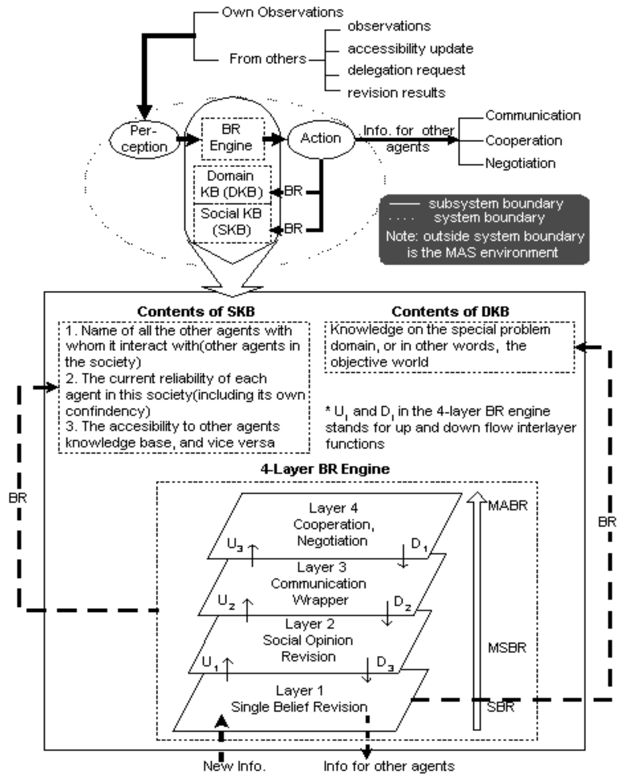

Figure 2: A Belief Revision Agent in MASs

For each layer, the input information is first checked by the *up flow* inter-layer function(the U_is in Figure 2). If there is no need to go up, then the intra-layer function is invoked to do the required processing. After that, the *down flow* function(the D_is in Figure 2) passes the processing results to the lower layer. Otherwise, the corresponding U_i passes the information handle to the upper layer and the same procedures are repeated in the layer above it. We take the inter-/intra layer functions of layer 1 for instance to describe the information flow in the BR Engine.

Layer 1 is the bottom layer and its lower layer is the environment. The environment includes the physical platform an agent lives in and all the other agents that it interacts with. Therefore, the input to layer 1 is either from its own observations or from other agents. Other agents could send it at least four types of information in the context of multi-agent belief revision, namely, their observations, the accessibility update of their knowledge bases, their revision result in a round of group/society wide revision or the delegation request. On taking the input from the agent perception system, U_1 first check whether the agent itself observes the new information ϕ:

- If yes, then it make decisions on whether to expose ϕ to other agents(e.g. checks the consistency of ϕ

[2]The agent society is defined from an $agent_i$'s point of view, it is populated with all the agents that it interacts with, including itself. Therefore, from different agent's viewpoint, the constitution of a society might be different. There are several ways for a agent to be known by the others, such as (i) register to an agent platform's yellow page as described by FIPA, (ii) broadcast its existence, or (iii) directly register to a particular agent's address book to join that agent's society.

	Intra-layer Function	Layer Interface			
				Input from a layer …	Output to a layer …
Layer 1	• Select revision strategies. • SBR using single/multiple sentence(s)	lower	own	Observations	• Own revision result after wrapping with wrapper passed from layer 3. • Information passed from layer 2.
			others	• Observations • Accessibility update • Revision result • Delegation request	
		upper		• New information credibility • Request on sending out information • The social opinion on the intended recipient(s)	• The handle of new information
Layer 2	• Update agents' reliability(incl. Self-confidence) • Determine SKS based on others' accessibility to its KB • Update its own Accessibility to other KBs on receiving announcement	lower		Same as output from layer 1 to upper layer	Same as the input of layer 1 from upper layer
		upper		• Communication wrapper (incl. Ontology, ACL*, communication mode, channel and etc.) with respect to the communication content • The intended recipient(s)	Communication desire generated if the new information is • revision results • delegation requests or after • checking the consistency of observations against current belief according to *knowledge Grade* • finding out consistency or inconsistency of different agent's KBs that it can access • updating accessibility of its own KB to others
Layer 3	• Determine the ontology • Generate speech act and formalize it using certain ACL* • Choose the communication mode(CLC** or COC***) • Specify the communication channel(i.e. to whom) • Determine whether to terminate a round of conversation on shared/common belief by waiting-for revision result	lower		Same as output from layer 2 to upper layer	Same as input of layer 2 from upper layer
		upper		• The type of delegation (or help) • The information on the conditions of starting or terminating negotiation	• A delegation request when intending to participate in teamwork or being aware of not being to do by itself • Offer of help
Layer 4	• Determine the type of delegation • Determine the type of help(incl. help level) • Choose the negotiation mechanism and strategy	lower		Same as output from layer 3 to upper layer	Same as input of layer 3 from upper layer
		uppe		N/A	N/A

*ACL: Agent Communication Language; **CLC: ConnectionLess Communication mode;
*** COC: Connection Oriented Communication mode

Table 1: Intra-layer functions and layer interface descriptions

against the various level of shared knowledge in the case of *knowledge migration*):

– if yes, pass ϕ to U_2.

– if no, invoke intra-layer function of layer 1. Selecting revision strategies according to defined *utility function* and revise the agent's own knowledge base.

- If no, pass the input to U_2.

For the rest layers please refer to table 1.

It is important to notice that the communication wrapper layer has no direct interactions with the environment. It only provides a wrapper consisting of an ontology, an ACL, a communication mode, etc for the corresponding information but does not send it to the environment directly. Interaction with the environment is left to layer 1. Each lower layer constantly provides basic functions to the layer above it. A layer can perform an action only through the interactions with all the layers under it.

This architecture allows us to achieve the flexibility of modeling various levels of BR within one framework. For example, the bottom layer could be detached to carry out Single BR; the bottom two layers could be separated to perform BR using information from multiple sources; similarly the bottom three layers can perform multi-agent belief revision with communication.

Shared Knowledge Structure(SKS)

Various types of heterogeneities which might affect multi-agent belief revision have been identified in Part I(Liu & Williams 1999). It is shown that a sophisticated knowledge structure is essential to capture the *social heterogeneity* so as to enable agents to share knowledge and preserve privacy. The major motivation of the development of a shared knowledge structure(SKS) is to find an automatic way of managing the accessibility of agent knowledge bases. Since an intelligent agent is considered to be proactive, it also should be proactive in the sense of dynamically arranging the face of its KB to different viewers.

In this paper, *information*, *knowledge*, *belief* are used interchangeably, except we sometimes distinguish knowledge from belief by emphasizing that flat knowledge takes the format of a sentence or a formula only, while belief is a piece of graded knowledge and consists of flat knowledge and a degree of belief.

Knowledge Classification via. Accessibility

The *accessible knowledge*(K_{Acc}) is defined as a subset of an agent's KB which is open to anonymous viewers or certain viewers specified by the agent itself. K_{Acc} could take various forms, such as a website open to information retrieval agents, or résumés which might be submitted to a specific job supplier agent. Based on K_{Acc}, the agent's KB can be constructed as described in this section.

¿From $agent_i$'s point of view, the *agent society* is populated by the agents that it interacts with. Let us assume that there are n agents in the society including $agent_i$. We denote the KB of $agent_i$ to be $K(i)$. It is composed of *m-accessible knowledge* and *m-semiprivate knowledge* as shown in Figure 3 and defined below:

1. *m-Accessible Knowledge:* Accessible knowledge for $agent_j$, denoted by $K_{Acc}(i,j)$, represents the set of formulae that $agent_i$ determines to open to $agent_j$. Since $agent_i$ can fully access its own knowledge base, when $i=j$, $K_{Acc}(i,i)$ is actually just $K(i)$. To avoid any confusion that might arise, $K(i)$ is adopted. Therefore, in the following sections, $K_{Acc}(i,j)$ only stands for the cases when $j \neq i$. As the subset of $K_{Acc}(i,j)$ may also be accessible to agents other than i and j, it is necessary to define strictly *2-accessible knowledge* ($K_{2-Acc}(i,j)$) and *accessible knowledge for 2 agents* ($K_{Acc}(i,j)$). Following is the definition:

$$K_{2-Acc}(i,j) = K_{Acc}(i,j) \cap \overline{(\cup(\cdots, K_{Acc}(i,k), \cdots))}$$

where $k \neq i$ and $k \neq j$. Let l be the cardinality of the set $\{K_{Acc}(i,j), \cdots, K_{Acc}(i,k), \cdots\}$, the set operation results in a strictly 2-Accessible knowledge set iff $l = C_{n-1}^{2-1} = n-1$, where C_s^r is the binomial coefficients, $\binom{s}{r} = \frac{s!}{(s-r)!r!}$,

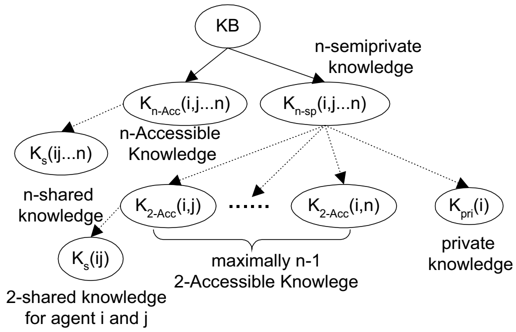

Figure 3: Shared Knowledge Structure

n is the population of the society. Therefore, the set of strictly 2-accessible knowledge $K_{2-Acc}(i,j)$ is the subset of $agent_i$'s knowledge that is only accessible from $agent_j$ but not any other agent in the society.

Given that m is the cardinality of the set $\{i,j,\cdots,k\}$, $1 < m \leq n$, the accessible knowledge for m agents $K_{Acc}(i,j\cdots k)$ is defined as follows(m is called the *degree of accessibility*):

$$K_{Acc}(i,j\cdots k) = \cap(K_{Acc}(i,j),\cdots,K_{Acc}(i,k))$$

Similar to the definition of $K_{2-Acc}(i,j)$, the *m-accessible knowledge*[3] $K_{m-Acc}(i,j\cdots k)$ could be defined based on $K_{Acc}(i,j\cdots k)$

$$\begin{aligned}&K_{m-Acc}(i,j\cdots k)\\&= K_{Acc}(i,j\cdots k) \cap \overline{(\cup(\cdots,K_{Acc}(i,p\cdots q),\cdots))}\end{aligned}$$

$K_{Acc}(i,p\cdots q)$ has the same degree of accessibility as $K_{Acc}(i,j\cdots k)$, but set $\{j,\cdots,k\}$ has no intersection with set $\{p,\cdots,q\}$. Let l be the cardinality of the set $\{K_{Acc}(i,j\cdots k),\cdots,K_{Acc}(i,p\cdots q),\cdots\}$, the right hand side of the set operation results in a strictly m-Accessible knowledge set iff $l = C_{n-1}^{m-1}$, n is the population of the society.

The bottom up m-accessible knowledge classification results in a tree structure, which is displayed in Figure 3, where $m = 1,\cdots,n$.

2. *m-Semiprivate/Private Knowledge:* As the complement of $K_{m-Acc}(i,j\cdots k)$, the m-semiprivate knowledge $K_{m-sp}(i,j\cdots k)$ is private towards $agent_j,\cdots,$ $agent_k$, which is accessible by $agent_i$ itself and agents other than $j,\cdots,k$. The following relation holds:

$$\begin{aligned}&\cup(\underbrace{\cdots K_{m-Acc}(i,j\cdots k)\cdots}_{C_{n-1}^{m-1}}) \cap K_{m-sp}(i,j\cdots k)\\&= K_{(m+1)-sp}(i,j\cdots k+1)\end{aligned}$$

When $m = 1$, the knowledge is only accessible by agent i itself and not by any one else, it becomes the *private knowledge* of $agent_i$, denoted by $K_{pri}(i)$.

[3] $m-accessible$ means strictly accessed by m agents. While the knowledge accessible by m agents means this knowledge is accessible by any subset of the m agents.

3. *KB of $agent_i$:* $K(i)$ is the knowledge base of $agent_i$. Within this KB, the agent can assign various level of accessibility to various agents in the society and reserve the rest of the KB as its private knowledge $K_{pri}(i)$.

$$\begin{aligned}&K(i)\\&= \cup(K_{pri}(i), \underbrace{K_{2-Acc}(i,j),\cdots,}_{C_{n-1}^{2-1}} \underbrace{K_{m-Acc}(i,j\cdots k),\cdots,}_{C_{n-1}^{m-1}}\\&\qquad \underbrace{K_{n-Acc}(i,j\cdots n))}_{1}\\&= \cup(K_{pri}(i), \underbrace{K_{Acc}(i,j),\cdots,K_{Acc}(i,k)\cdots,K_{Acc}(i,n))}_{n-1}\end{aligned}$$

where $j \neq i$ and $j = 1,\cdots n$ for an n agent society.

Applying the definitions above to the other agent's knowledge bases, $K(i)$ can be further classified:

1. *Mutually Accessible Knowledge:* It is the set of formulae that agents $i,j,\cdots,k$ all believe but may not know whether everyone knows that everyone believes that ... and may not know to what degree. 2-mutually accessible knowledge $K_{AccMutual}(ij)$ is a special case when both agents believe and know each other knows that each of them believes, which is defined as:

$$K_{AccMutual}(ij) = \cap(K_{acc}(i,j),K_{acc}(j,i))$$

The *the ideal m-mutually Accessible Knowledge* defined in the similar way is not realizable. Since the knowledge is classified from $agent_i$'s viewpoint, it is impossible for i to know about the knowledge classification of other agents.

$$\begin{aligned}&K_{AccMutualIdeal}(ij\cdots k)\\&= \cap(\underbrace{K_{Acc}(i,j\cdots k),\cdots,K_{Acc}(j,i\cdots k),K_{Acc}(k,i\cdots j)}_{\frac{m!}{C_m^{m-1}}=m}\end{aligned}$$

where, m is the cardinality of set $\{i,j,\cdots,k\}$.

Therefore, only super agents who can access every agent's KB can do the above evaluation. For agent i, it could only reason about the possible m-mutually accessible knowledge at a coarse level:

$$\begin{aligned}&K_{AccMutualPossi}(ij\cdots k)\\&= \cap(K_{AccMutual}(ij),\cdots,K_{AccMutual}(ik))\\&\supseteq K_{AccMutualIdeal}(ij\cdots k)\end{aligned}$$

but

$$\begin{aligned}&K_{AccMutual}(ij)\\&= K_{AccMutualPossi}(ij) = K_{AccMutualIdeal}(ij)\end{aligned}$$

2. *m-Shared Beliefs:* $K_s(ij)$ represent the *2-shared knowledge* of agent $i, and\ j$. It describe a package of knowledge that agent i and j both believe and both of them know that each other believes it to *an agreed degree*. $K_s(ij)$ is defined as:

$$K_s(ij) = (*)K_{AccMutual}(ij)$$

where $*$ stands for the various kinds of operator that can generate shared knowledge from the mutually accessible knowledge. One simple candidate is to simply take the mutually accessible knowledge as shared knowledge and derive a unanimous ranking on it. Alternatively it could be used to initiate conversations among the agents to select an agreed subset.

m-shared beliefs is defined as

$$K_s(ij \cdots k) = (*) \cap (K_{sij}, \cdots, K_{sik})$$

where $*$ has the same meaning as defined above. The discussion for possible m-mutually accessible knowledge suggests that $*$ in this case should enable $agent_i$ to send out queries to other agents. This is to determine the accessibility of $K_{AccMutualPossi}(ij \cdots k)$ from the other agents viewpoint. If all the agents faithfully answer queries whenever required, $K_{AccMutualIdeal}(ij \cdots k)$ would be the final result of such queries.

3. *Common Beliefs:* K_c is the n-shared beliefs, which is a special case of shared beliefs when the set of formulae is believed by the whole society with an agreed degree(i.e. n-shared knowledge), and everyone in the society knows that everyone believes it at a certain degree and $\cdots$ and everyone knows that $\cdots$ that everyone believes it to a certain degree.

The mutually accessible knowledge is stored separately in each agent's local KB, the rank of the same formula in $K_{AccMutual}(ij \cdots k)$ does not necessarily have to be the same according to individual agent's judgment. But when the agents decide to share some knowledge, it is necessary for them to agree on the rank of the shared knowledge. Thus the rest of each agent's KB needs to be revised to accommodate the agreed degree. By the way, shared/common beliefs could be extracted from each agent's KB and stored in a common location if everyone agrees to do so. This would be beneficial in terms of saving space and it would enhance the robustness of the dynamics of shared information.

Relationship with Other Knowledge Structures

Essentially, two types of knowledge structures for modeling common/shared knowledge exist in the literature, (i) the labeled tree of mutual belief by Van de Meyden(Meyden 1994) based on *possible worlds* semantics of modeling knowledge and belief in MASs(Fagin *et al.* 1995); and (ii) the knowledge in a Shared Domain (SD) and a Private Domain (PD) by Kfir-dahav and Tennenholtz(Kfir-dahav & Tennenholtz 1996).

In the case of mutual belief(Meyden 1994), the accessibility relations are defined on the possible worlds according to the agents current beliefs. A labeled tree formally represents this, which is a mixture of domain and social knowledge. Each agent is able to reason uniformly because the whole tree structure is visible to everyone. It is inflexible to some extent because it does not support private knowledge.

Tidhar et al.(Tidhar, Sonenberg, & Rao 1998) define *team knowledge* as "the team knows only what is known to every subteam". Such team knowledge is actually the intersection of subteams' knowledge base, which is similar as *SD* defined in (Kfir-dahav & Tennenholtz 1996). It is shown in Part I(Liu & Williams 1999) such a classification is not capable of modeling *truely* private knowledge. It implicitly presumes a super agent who can access all the agents' KBs, draw the conclusions from the shared knowledge and impose that on each individual. Imposing team knowledge top down from a team leader is only one way of achieving team knowledge. The other way is to derive concensus (via. negotiation or argumentation etc.) from the team members. The latter way is not supported by the structures in (Kfir-dahav & Tennenholtz 1996)(Tidhar, Sonenberg, & Rao 1998).

Intuitively, the agents knowing the same thing does not necessarily imply either the awareness of, or the sharing of intentions. While agents sharing some knowledge means the shared knowledge must be known to each agent involved, i.e. Shared $\overset{\Rightarrow}{\not\Leftarrow}$ Known. Previous research reported in the literature fails to capture such intuition. However our SKS is successful in achieving it by authorizing agents to organize their own knowledge bases. The accessibility as defined in section is totally determined by an individual agent's personal stance. It is considered as an attribute of a certain subset of the agent KB regardless of how the agent is going to reason on it. Since there is no need to reveal all of its KB, true privacy is maintained when cooperating. The operator "(*)"is suggested to derive shared belief from a rational set of *mutually accessible knowledge* without the aid of super agents. However, this knowledge structure does not preclude the possibility of the existence of a super agent. This could be done by simply forcing all the other agents to register its platform to give full accessibility to the super agent. The super agent, on the other hand, would allow only limited or no accessibility to others.

The SKS as proposed is capable of simulating previous structures by rearranging the accessibility relations. For example, if every agent opens its domain knowledge base to all the others, the *n-possible mutually accessible knowledge* become *n-ideal mutually accessible knowledge.* The n-shared knowledge achieved is just the knowledge in *SD*. The complement is left for *PD*. If every agent fully opens both its SKB and DKB, the accessibility relation become *transitive.* Therefore, SKS is able to capture the semantics of mutual belief(Meyden 1994), and in this case, $K_{pri}(i)$ is empty.

Knowledge Migration

Inconsistency Principle

In this framework, the consistency of an individual's KB with the common/shared beliefs is our primary concern. This is achieved using standard single agent BR techniques. So local consistency is still a prerequisite of each agent, but the global consistency of the society is not required either during or after the process of BR. Hence, inconsistency across the KBs of the society is permitted. This is called *inconsistency principle*. This is different from the *Liberal Belief Revision Policy*(Dragoni, Guorgini, & Baffetti 1997) where the final DBR goal is to achieve global consistency.

Observation: If formula p is in consistently believed in the agent society, then neither $p \in K_c$ nor $\neg p \in K_c$

This observation also holds for the shared belief bases. That is, among the agents who share the belief, neither it or its negation are derivable from the shared belief base. And it becomes one of the key postulates that should be satisfied during multi-agent belief revision. Therefore, as long as a formula is not derivable from the common belief base, it could be inconsistently believed across the society or across several groups; similarly, as long as a formula is not in a shared belief base of a certain group, it could be held inconsistently *at least* across the agents in this group or across several subgroups.

The major goal of BR based on the proposed SKS and the *inconsistency principle* is to maintain the consistency of $K(i)$ with K_c, and with shared belief base $K_s(ij \cdots k)$.

Common/shared beliefs are distinct from mutually accessible knowledge by the fact that the rank in $K_c/K_s(ij \cdots k)$ should be respected, while different ranks could be assigned to the same mutually accessible knowledge by different agents. Therefore, maintaining the consistency of $K(i)$ with K_c means aligning the rank of common beliefs in every agents KB when new information is accepted by the society. That is, the revision process to revise $K(i)$ by a set of sentences with fixed ranks (in $K_c/K_s(ij \cdots k)$) and also the new information, while the rank of new information and sentences in $K(i)$ are changeable, those in K_c are not[4]. Revision regarding the shared knowledge $K_s(ij \cdots k)$ uses a similar process but within a smaller groups of agents. Communication and negotiation might be necessary to reach the concensus on ranking of $K_c/K_s(ij \cdots k)$.

Consequently, belief dynamics considered here are more sophisticated than that is required for single agent revision and could be extended to *knowledge migration* in terms of *knowledge grade*, which is discussed in the next section.

Knowledge Grade and Migration

Based on the SKS, an agent's KB could be ordered according to the degree to which beliefs are shared:

Common Beliefs	$\uparrow$ High
$(n, \cdots, 2)$-Shared Beliefs	|
$(n, \cdots, 2)$-AccNotShared Beliefs	|
Private Beliefs	| Low

where n is the agent population. Common beliefs have the highest grade of sharing and private beliefs have the lowest.

According to the *inconsistency principle*, inconsistency is permitted across the agent society as long as the inconsistent knowledge is not in K_c. Therefore, if the new information $\neg\phi$ contradicts the common belief, instead of giving up ϕ and expanding K_c by $\neg\phi$ in the traditional way, the following procedure could be taken to incorporate $\neg\phi$ as well as letting some agents whose background knowledge supports ϕ still keep it:

- The observer agent *broadcasts* the observation and waits for response. If all the agents agree to do nothing to K_c, then ignore $\neg\phi$ and keep everything unchanged. If all agree to accept $\neg\phi$, perform *mirror revision* on K_c with respect to ϕ and derive consensus on the result $K_c \dot{+} \phi$. Otherwise, perform *mirror contraction* on K_c with respect to ϕ and deriving consensus on the result $K_c - \phi$ and then go to the next procedure.
- There is a maximum of $n-1$ groups of agents who share their own $n-1$ *shared beliefs*. Check within each group whether they wish to accept, say, $\neg\phi$. If yes, revise the corresponding shared belief base with $\neg\phi$. If no, check whether to accept ϕ. For the group who won't take either one, the process repeats through the smaller subgroups until they are broken down into the individual agent.

In this procedure, an important phase is that the original believed common belief ϕ has been migrated down either to the some groups' shared belief bases $K_s(ij \cdots k)$ or to someone's $K_{pri}(i)$. In the sense of *knowledge grade*, ϕ has been *degraded* during the revision.

On the other hand, information could also be *upgraded* from a lower grade to a higher one. At first, the revision on accessible but not shared knowledge results in the accessibility update. Then, when an agent updates its accessibility or receives other agents' accessibility update announcement, it will reevaluate the knowledge in its KB and the accessible knowledge in others' KBs. Following the procedure of defining shared/common belief in section , m-shared belief can be generated from the lower grade m-accessible knowledge. In fact, the process of deriving shared/common knowledge is just the process of *knowledge upgrade*.

The process of knowledge upgrade and degrade is called *knowledge migration*. Scenarios of *knowledge migration* may vary according to the application. In general, knowledge migration is the process of reshaping the knowledge structure triggered by the new information.

Conclusions and Future Work

The multi-agent belief revision framework developed herein provides the necessary infrastructure to support several levels of BR e.g. Single BR, BR using information from multiple sources and full-blown team-based

[4]This is called *mirror revision*, and correspondingly, *mirror contraction*.

multi-agent BR. Thus heterogenous BR agents can be modeled under one framework. To tackle the social heterogeneities in multi-agent belief revision, a SKS is proposed to enable the sharing of knowledge as well as maintaining private knowledge. Since the SKS is applicable to both the SKB and the DKB, it is capable of modeling "legacy"knowledge structures. Based on the *inconsistency principle*, and the *knowledge grade* defined by SKS, the process of multi-agent belief revision becomes the process of knowledge migration.

Many other features of this framework is still under investigation and need future work. For example, on the accessibility relation, agents could choose whether to reveal the rank on beliefs or not. This would allow modeling various social attitudes[5] that might be taken during cooperation and teamwork.

Applying the SKS to the SKB will enable agents to propagate their social opinion e.g. its evaluation of other agent's trustworthiness, its accessibility to others' KB and so on. One interesting benefit is that using propagated social opinion, agents could have a natural way of evaluating other agent's reliability. The knowledge migration process can be made more sophisticated when taking the communication mode (i.e connectionless/connection-oriented) and protocols(i.e. broadcast, multicast and peer to peer) into consideration.

References

Castelfranchi, C.; Rosis, F. D.; Falcone, R.; and Pizzutilo, S. 1998. Personality traits and social attitudes in multiagent cooperation. *Applied Artificial Intelligence* 12:649–675.

Davis, R., and Smith, R. 1983. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence* 20:63–109.

Dragoni, A., and Guorgini, P. 1999. Revising beliefs received form multiple sources. In Williams, M. A., and Rott, H., eds., *Frontiers of Belief Revision*.

Dragoni, A.; Guorgini, P.; and Baffetti, M. 1997. Distributed belief revision vs. belief revision in a multiagent environment: First result of a simulation experiment. In *Multi-Agent Rationality, Proc. of MAAMAW'97*, LNAI, 45–62. Ronneby, Sweden: Springer-Verlag.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge*. London: The MIT Press, 1 edition.

Kfir-dahav, N. E., and Tennenholtz, M. 1996. Mutiagent belief revision. In *6th Conf. on Theoretical Aspects to Rationality and Knowledge*, 175–194.

Liu, W., and Williams, M. A. 1999. A framework for multi-agent belief revision (part i: The role of ontology). In Foo, N., ed., *12th Australian Joint Conference on Artificial Intelligence*, Lecture Notes in Artificial Intelligence. Sydney, Australia: Springer-Verlag.

Meyden, R. v. d. 1994. Mutual belief revision. In *the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR'94)*.

Müller, J. P. 1998. Architectures and applications of intelligent agents: A survey. *The Knowledge Engineering Review* 13(4):353–380.

SATEN. http://u2.newcastle.edu.au/webworld/saten/ (hint: use netscape 4.5 or above with java 1.1.1 plugin).

Tidhar, G.; Sonenberg, E. A.; and Rao, A. S. 1998. On team knowledge and common knowledge. In Demazeau, Y., ed., *Third International Conference on Multi-Agent Systems*, 301–308. Paris, France: IEEE Computer Society.

Williams, M. A. 1998. Applications of belief revision. In Freitag, B., ed., *Transactions and Change in Logic Databases*, volume LNCS 1472. Springer Verlag. 15–47.

Wooldridge, M. 1999. Intelligent agents. In Weiss, G., ed., *Multiagent Systems*. The MIT Press.

[5]including benevolent, malicious(Liu & Williams 1999) and the helping attitudes in (Castelfranchi *et al.* 1998).