Logique & Analyse **133-134** (1991), 121-132

A SUBJECTIVE LOGIC OF KNOWLEDGE

Ph. LEJOLY

1. *Introduction*

A robot (Rob1) moves slowly in a room. His batteries are nearly down. He is desperately seeking an electrical outlet, but there is none in sight in this room. Rob1 knows that there must be a plug either in the room to the left, or in the room to the right, but he only has enough power to visit *one* more room. Without any further information, Rob1 sadly decides to use his random number generator, when he sees his colleague Rob2 coming out of the room to the left. "Well", Rob1 thinks, "I don't know whether there is an outlet in that room, but I know that Rob2 knows, since he has been there. So, I'll ask him." After Rob2's answer, Rob1 is able to decide where to go.

This kind of reasoning, taking into account the "mental states" of agents —in order to make a decision, Rob1 uses not only what he knows about the environment, but also what he knows about Rob2's knowledge— can be formalized using the *logic of knowledge*, or *epistemic logic* ([Hin62, HM85, Hal86]).

One approach is to axiomatize the situation of interest as a set of epistemic logic formulas and use a deductive system to infer conclusions (agents' actions for instance). This approach suffers from the usual weaknesses of deductive systems, even to a greater degree, since the logic of knowledge is more difficult to handle than classical logic.

An alternative is to represent the situation of interest by a structure and to evaluate formulas on this structure (for instance, agents' actions can be determined by checking if the antecedent of some rule is satisfied on the structure). This structure will be of the same type as a model of the logic and evaluation will follow the semantic rules of the logic. This approach is known under the name of *model-checking* ([HV91]) since it amounts to checking that a structure is a *model* of a given formula.

In many applications of the logic of knowledge, the notion of knowledge considered is an *external one*: it is a tool used by the system designer to express and verify certain properties of the system ([HM84, FHV86, RK86, FI86]). But a component of the system cannot use what it "knows", only the designer performs epistemic reasoning.

Rather, we would like to take the point of view of the agent, i. e. an *internal* or *subjective* approach. In this approach, *the agent* performs epistemic reasoning. This amounts to providing each agent with a model *of his own knowledge* which *he* will use to evaluate formulas, and which *he* will update when he gets new information. There is no structure describing the global situation any more, or only as a conceptual tool inaccessible to a particular agent. So, an agent would start with an empty model and would update it as he "perceives" facts, or as he communicates with other agents. To determine if a given formula is true, an agent evaluates it on his model. This allows us to give a realistic model of interacting agents and of their inference capabilities. For instance our agents have only a partial view of the situation. They are aware of a limited subset of primitive propositions and their current model is a representation of the finite amount of information they have gained thus far.

We use a formalism based on the notion of knowledge structures ([FHV84, FV85, FHV91]), rather than on the classical *Kripke structures* ([Kri63, HM85, Hal86]). A problem is that these knowledge structures are infinite, and thus cannot be manipulat1 3 4o provide each agent with a *finite representation* of his model, along with a way to evaluate formulas on it. This is done by introducing circularity, or self-reference, in the definition of a model. Among other things, this allows a finite representation of the common knowledge of a group of agents, and the evaluation in a finite number of steps of common knowledge formulas. Furthermore, we present algorithms for updating the model of an agent when his information changes due to communication.

2. *The Model*

We can view an agent's knowledge as consisting of two parts:

1. knowledge about the environment (level-0 knowledge);
2. knowledge about the knowledge of the other agents (*meta-knowledge*).

This leads us very naturally to define the model of an agent as a set of *worlds*, each one describing a state of affairs considered possible by the agent and having a two-fold structure, one part representing level-0 knowledge and the other part representing meta-knowledge.

As usual, an agent *knows* a formula if it is true on each world he considers possible, i. e. if it is true on his model.

In the propositional case, level-0 knowledge will be represented by a truth value assignment to primitive propositions. However, we would like to take care of the fact that an agent has only a limited awareness. For example, let's imagine that agent i is a five year old boy and p is the proposition

The complexity class P is identical to the complexity class NP.

It is quite conceivable that our boy is absolutely not aware of proposition p. For him, the problem of attributing a truth value to p does not even exist! An agent's model will therefore contain information only about those primitive propositions the agent is aware of.

The meta-knowledge part of a world will be represented by associating to each agent a model, i. e. a set of worlds.

More formally, let A be a finite set of agents and P a finite set of primitive propositions.

The set of worlds W is defined as follows:

$$W = \langle L_0, L_K \rangle \in \mathrm{W} \text{ iff}$$

- $L_0 : S \subseteq P \rightarrow \{T,F\}$
- $L_K : A \rightarrow 2^{\mathrm{W}}$

This means that an object $W = \langle L_0, L_K \rangle$ belongs to the set of worlds if and only if it has the desired structure, i.e. if L_0 is a partial truth value assignment and L_K associates to each agent a set of worlds.

For each agent i, the model M_i of i is a set of worlds, i.e. a subset of W.

Given the models M_i of the different agents, we can define the *global model*

$$M_G = \{\langle G_0, G_K \rangle\}$$

where

- G_0 is a *complete* truth value assignment, representing the "real" environment, i.e. $G_0 : P \rightarrow \{T,F\}$
- $G_K(i) = M_i$

So, the global model is formed by grouping together the individual models plus the real description of the environment. Note that this global model is only a conceptual tool, inaccessible to a particular agent, who has only his own (incomplete) model to work with.

We can already make two remarks about the previous definitions:

- The model of an agent has a *finite representation*. This allows it to be manipulated by the agent, either to evaluate formulas on it, or to update it with new information.

- Circularity (self-reference) is allowed in the definition of a world. For instance, we can have $W = \langle L_0, L_K \rangle$ with W itself belonging to $L_K(i)$: W is then a part of its own definition. As we will see, this makes it possible to model infinitary notions like introspection and common or joint knowledge within this finite framework.

Structural constraint

In order to maintain coherence between successive levels of knowledge, we impose the following structural constraint on worlds:

(K1) Let $W = \langle L_0, L_K \rangle \in \mathsf{W}$. Then, $\forall W' = \langle L'_0, L'_K \rangle \in L_K(i)$, $L'_K(i) = L_K(i)$ for all agent i.

This means that the worlds considered possible by an agent at the next level of knowledge (i.e. the worlds he considers possible when in a world $\langle L'_0, L'_K \rangle$ considered possible at the present level of knowledge) are the same as those considered possible at the present level of knowledge.

As we will see, the constraint (K1) implies that, at agent i 's level, for any formula φ, φ is equivalent to $\mathsf{K}_i\varphi$.

3. *Evaluation of formulas*

Let M_i be the model of agent i and φ a formula of a propositional modal logic. As stated above, M_i satisfies φ if and only if every world in M_i satisfies φ. But remember that we want to take the point of view of the agent. This means that we want to define how *an agent* evaluates formulas on his model. So, rather than defining satisfiability of a formula on a world in general, we define satisfiability of a formula on a world *relatively to* an agent. So, $W, i \models \varphi$ means "the formula φ is satisfied on the world W relatively to agent i ".

Satisfiability on an agent's model now becomes:

$$M_i \models \varphi \textit{ iff } W, i \models \varphi, \forall W \in M_i$$

We define satisfiability on a world $\langle L_0, L_K \rangle$, relatively to an agent i as follows:

$$W, i \models p \textit{ iff } L_0(p) = T, p \in P \quad (1)$$
$$W, i \models \neg p \textit{ iff } L_0(p) = F, p \in P \quad (2)$$
$$W, i \models \varphi \wedge \psi \textit{ iff } W, i \models \varphi \text{ and } W,i \models \psi \quad (3)$$
$$W, i \models \varphi \vee \psi \textit{ iff } W, i \models \varphi \text{ or } W,i \models \psi \quad (4)$$
$$W, i \models \neg(\varphi \wedge \psi) \textit{ iff } W, i \models \neg\varphi \text{ or } W,i \models \neg\psi \quad (5)$$
$$W, i \models \neg(\varphi \vee \psi) \textit{ iff } W, i \models \neg\varphi \text{ and } W,i \models \neg\psi \quad (6)$$
$$W, i \models \mathsf{K}_j\varphi \textit{ iff } W', j \models \varphi \ \forall W' \in L_K(j) \quad (7)$$
$$W, i \models \neg\mathsf{K}_i\varphi \textit{ iff } \exists W' \in L_K(i) \text{ s.t. } W', i \models \neg\varphi \quad (8)$$
$$W, i \models \neg\mathsf{K}_j\varphi \textit{ iff } W, i \not\models \mathsf{K}_j\varphi, j \neq i \quad (9)$$

The definitions (2), (5), (6) and (8) take care of the partiality of an agent's awareness (we can no longer define the satisfiability of $\neg\varphi$ in terms of the non-satisfiability of φ). Note that, given a formula φ, it is possible that neither φ, nor $\neg\varphi$ are satisfied on a world.

Definition (7) expresses the fact that, in order to evaluate the formula "agent j knows φ" on a world W, agent i has to shift his point of view: he evaluates φ on the worlds associated to j , but he does so *relatively to* j. In other words, agent i evaluates φ on the model associated to j in W *as if* he were agent j.

Of course, satisfiability on the global model $M_G = \{\langle G_0, G_K \rangle\}$ is defined straightforwardly as follows:

$$M_G \models p \textit{ iff } G_0(p) = T, p \in P \quad (10)$$
$$M_G \models \neg\varphi \textit{ iff } M_G \not\models \varphi \quad (11)$$
$$M_G \models \varphi \wedge \psi \textit{ iff } M_G \models \varphi \text{ and } M_G \models \psi \quad (12)$$
$$M_G \models \varphi \vee \psi \textit{ iff } M_G \models \varphi \text{ or } M_G \models \psi \quad (13)$$
$$M_G \models \mathsf{K}_i\varphi \textit{ iff } G_K(i) \models \varphi \quad (14)$$

Remember that, by construction, $G_K(i) = M_i$, so (14) is equivalent to:

$$M_G \models \mathsf{K}_i\varphi \textit{ iff } M_i \models \varphi \tag{15}$$

If the constraint (K1) holds, it is easy to see that the following formulas are agent-valid (i.e. true in every agent's model):

$$\varphi \Rightarrow \mathsf{K}_i\varphi \tag{16}$$
$$\mathsf{K}_i\varphi \Rightarrow \varphi \tag{17}$$

As particular instances of (16),

$$\mathsf{K}_i\varphi \Rightarrow \mathsf{K}_i\mathsf{K}_i\varphi \tag{18}$$
$$\neg\mathsf{K}_i\varphi \Rightarrow \mathsf{K}_i\neg\mathsf{K}_i\varphi \tag{19}$$

are also agent-valid.

Formulas (16) and (17) express that, taking the point of view of an agent, "φ is true" is equivalent to "I know that φ is true". From this follow the introspective capabilities of the agents expressed by (18) and (19). Positive introspection (18) seems very natural, whereas negative introspection (19) is often rejected as unrealistic. But this is not the case here: (19) is as natural as (18). Indeed, what (19) (evaluated by agent i on his model) expresses is: "if I, agent i, do not know φ, then I know that I do not know φ". How could it be otherwise: for the antecedent to be true, agent i must have performed the process of checking whether he knows φ and have arrived at the answer "no", thus *knowing* this answer. Of course, negative introspection *is* unrealistic at the global level but, there, it does not hold anymore. For instance, suppose that agent i is not aware of primitive proposition p. In this case, both $\neg\mathsf{K}_i p$ and $\neg\mathsf{K}_i\neg p$ are satisfied at the global level, but $\neg\mathsf{K}_i p$ is not satisfied on agent i 's model. So $\mathsf{K}_i\neg\mathsf{K}_i p$is not satisfied at the global level, although $\neg\mathsf{K}_i p$ is.

Using the definition (15) above, we can see that the following formula is valid at the global level:

$$[\mathsf{K}_i\varphi \ \wedge \ \mathsf{K}_i(\varphi \Rightarrow \psi)] \Rightarrow \mathsf{K}_i\psi \tag{20}$$

Formula (20) expresses the closure of knowledge under implication. In the "classical" theory, this, combined with the fact that all propositional tautologies are known, gives rise to what is called *logical omniscience*: the agents know all the consequences of what they know. Here, only a limited form of logical omniscience remains. An agent does not know *all* propositional tautologies but only those composed of primitive propositions he is aware of. So, he is logically omniscient in his awareness universe only.

Moreover we can remark that "the well-known *logical omniscience* problem does not present the same difficulties in the model-checking approach as it does in the theorem-proving approach" ([HV91], p. 1). An agent knows all the tautologies composed of primitive propositions he is aware of, but he does not know whether a particular formula is a tautology or not. If asked to evaluate a valid formula, he will answer "true" since it will be true in his current model, but he is incapable of determining whether the formula is valid or serendipitously true.

3.1 *Example*

Suppose we have two agents and the model of agent 1 is the following:

$$
\begin{aligned}
M_1 &= \{W^{11}, W^{12}\}, \text{ with} \\
W^{11} &= \langle \{p, q, r\}; 1 \rightarrow M_1, 2 \rightarrow \{W^{21}\}\rangle \\
W^{12} &= \langle \{p, \neg q, r\}; 1 \rightarrow M_1, 2 \rightarrow \{W^{22}\}\rangle \\
W^{21} &= \langle \{q, r\}; 1 \rightarrow W^{0}\}, 2 \rightarrow \{W^{21}\}\rangle \\
W^{22} &= \langle \{\neg q, r\}; 1 \rightarrow W^{0}\}, 2 \rightarrow \{W^{22}\}\rangle \\
W^{0} &= \langle \varnothing; 1 \rightarrow \varnothing, 2 \rightarrow \varnothing\rangle
\end{aligned}
$$

This corresponds to the following situation: agent 1 considers two states of affairs as possible (W^{11} and W^{12}). In the first one, p, q and r are true and in the second one p and r are true whereas q is false. So, agent 1 knows that p and r are true, but does not know whether q is true or false, although he is aware of q. Moreover, in W^{11} agent 1 assigns as only possibility to agent 2 a world in which q (and r) is true. So in W^{11}, agent 2 knows that q is true. Similarly, in W^{12}, agent 2 knows that q is false. So, although he does not know whether q is true or false, agent 1 knows that agent 2 knows it.

We can check that the following propositions hold:

$$M_1 \models \neg \mathsf{K}_1 q \wedge \neg \mathsf{K}_1 \neg q \tag{21}$$

$$M_1 \models \mathsf{K}_2 q \vee \mathsf{K}_2 \neg q \tag{22}$$

$$M_1 \models \mathsf{K}_1 \neg \mathsf{K}_1 q \tag{23}$$

$$M_1 \not\models \mathsf{K}_2(p \vee \neg p) \tag{24}$$

$$M_1 \models \neg \mathsf{K}_2 p \tag{25}$$

$$M_1 \not\models \mathsf{K}_2 \neg \mathsf{K}_2 p \tag{26}$$

$$M_1 \not\models z \vee \neg z \tag{27}$$

Proposition (24) holds because, in the possible worlds assigned to agent 2 by agent 1 (W^{21} or W^{22}), p has no truth value. This expresses the fact that agent 1 considers that agent 2 is not aware of p.

To see why (26) holds, let us perform the evaluation of $\mathsf{K}_2\neg\mathsf{K}_2p$ on M_1 : $\mathsf{K}_2\neg\mathsf{K}_2p$ is satisfied on M_1 iff it is satisfied on every world in M_1, relatively to 1. Consider W^{11} : $W^{11}, 1 \models \mathsf{K}_2\neg\mathsf{K}_2p$ iff $W^{21}, 2, \models \neg\mathsf{K}_2p$ (take note of the shift of point of view). By (8), $W^{21}, 2, \models \neg\mathsf{K}_2p$ iff $\exists W \in \{W^{21}\}$ s.t. $W, 2 \models \neg p$. But $W^{21} \not\models \neg p$. So (26) holds.

4. *Common Knowledge*

Intuitively, a proposition φ is common (or joint) knowledge to a group G of agents if φ is true, and everyone in G knows that φ is true, and everyone in G knows that everyone in G knows that φ is true, and so on. Formally, this gives rise to the following *iterative* definition:

$$C_G\varphi \equiv \varphi \wedge (\bigwedge_{i\in G}\mathsf{K}_i\varphi) \wedge (\bigwedge_{i,j\in G}\mathsf{K}_i\mathsf{K}_j\varphi) \wedge \ldots$$

But this iterative definition of common knowledge, in terms of an infinite conjunction of formulas of increasing depth, does not lend itself to evaluation in finitely many steps: if agent 1 tries to evaluate $C_{\{1,2\}}p \equiv p \wedge (\bigwedge_{i\in\{1,2\}}\mathsf{K}_ip) \wedge (\bigwedge_{i,j\in\{1,2\}}\mathsf{K}_i\mathsf{K}_jp) \wedge \ldots$, he will of course evaluate each of the components of the conjunction to true, but he will go on forever.

We can however give an equivalent definition of common knowledge which fits nicely with our formalism. A fact φ is common knowledge to a group of agents means that the knowledge state of each agent in the group is such that φ holds in it and each agent in the group knows that this knowledge state holds. This leads to define $C_G\varphi$ as a *fixed point* ([Mos86, Bar88]):

$$C_G\varphi \equiv \varphi \wedge \bigwedge_{i\in G}\mathsf{K}_i(\varphi \wedge C_G\varphi)$$

This enables finite evaluation of common knowledge formulas. For instance, consider a world $W = \langle L_0, L_K\rangle$. In order to evaluate $C_{\{i,j\}}\varphi$ on W, we have first to evaluate φ, $\mathsf{K}_i\varphi$ and $\mathsf{K}_j\varphi$ on W, i.e. to evaluate φ on the worlds in $\{W\} \cup L_K(i) \cup L_K(j)$. Then we have to repeat the evaluation of $C_{\{i,j\}}\varphi$ on those worlds. Due to the particular structure of the worlds, this procedure will not go on forever: sooner or later, a loop will lead us to evaluate $C_{\{i,j\}}\varphi$ on worlds on which this evaluation is already in progress.

Formally, in order to evaluate $C_G\varphi$ on a world $W = \langle L_0, L_K\rangle$, we construct the sequence of sets of worlds:

$$S^0 = \{W\}$$
$$S^{n+1} = S^n \cup \{ \bigcup_{i \in G} L'_K(i), \forall \langle L'_0, L'_K\rangle \in S^n\}$$

until a fixed point is reached, i.e. $S^{n+1} = S^n$.
Then,

$$W \models C_G\varphi \textit{ iff } W' \models \varphi, \forall W' \in S^{n+1}$$

5. *Update*

We are now going to explain how an agent can update his model with new information. We consider the simplest framework, where we make the following hypotheses:

- The environment is *static* . By this we mean that the primitive propositions keep the same truth value over time.
- Agents are *honest* , i. e. they communicate only propositions they consider true.

So, at a given time, an agent has only partial knowledge of the global system composed of the environment plus the different agents. His knowledge will become more and more precise as he gets new information. This information can come either directly from the environment (the agent "perceives" something), or from another agent.

Now, suppose that the model of agent i is M_i. In order to update M_i with a proposition φ, i constructs a new model M'_i by "adding" φ to M_i. Note that, even in the simple framework we consider, an agent can receive information inconsistent with the set of formulas satisfied on his current model. For instance, suppose that agent 1 knows that agent 2 does not know whether p is true or false, i.e. $(\neg K_2 p \wedge \neg K_2 \neg p)$ is satisfied on M_1. Later on, agent 2 gets precise information about p, for example that p is true, and he communicates this to agent 1. So, agent 1 receives the message "p is true" from 2. He has then to update his model with $(K_2 p)$, which is inconsistent with $(\neg K_2 p \wedge \neg K_2 \neg p)$.

So, although the environment is static, the knowledge and meta-knowledge the agents have thereof is dynamic, and this leads to knowledge revision as above.

Our goal then is to set up a function *Add* which takes a set of worlds S, a formula φ and an agent i and returns the set S' obtained by adding to each world in S, relatively to agent i. We want *Add* to satisfy the following properties:

1. for all world $W' \in Add(S, \varphi, i)$, $W', i \models \varphi$,
2. S' is a "refinement" of S with the new information φ.

We give below the definition of Add for a primitive proposition p:

- $Add(S, p, i) = \bigcup_{\langle Lo, LK\rangle \in S} \{\langle L_0 + (\neg)p, L'_K\rangle\}$

  such that

  $$L'_K(i) = Add(L_K(i), (\neg)p, i)$$
  $$L'_K(j) = L_K(j), j \neq i$$

- $Add(S, \varphi \wedge \psi, i) = Add(Add(S, \varphi, i), \psi, i)$
- $Add(S, \varphi \vee \psi, i) =$
  $\bigcup_{M \in S} \{Add(\{M\}, \varphi \wedge \neg\psi, i) \cup Add(\{M\}, \neg\varphi \wedge \psi, i) \cup Add(\{M\}, \varphi \wedge \psi, i)$
- $Add(S, \mathsf{K}_j\varphi, i) = \bigcup_{\langle Lo, LK\rangle \in S} \{\langle L_0, L'_K\rangle\}$

  such that

  $$L'_K(i) = Add(L_K(i), \mathsf{K}_j\varphi, i)$$
  $$L'_K(j) = Add(L_K(j), \varphi, j)$$

Property : If the world W satisfies the constraint (K1), then $Add(\{W\}, \varphi, i)$ satisfies (K1).

6. *Conclusion and Comparison with Other Work*

Our work is an attempt to define an *internal*, or *subjective* notion of knowledge, by focusing on the agent rather than on the system. We set up a formalism which provides each agent with a model of his own. This model has a *finite representation*, allowing manipulation and modification by the agent, and finite evaluation of common or joint knowledge formulas. Compare this with the "classical" knowledge structures, where the satisfaction of these formulas needs the definition of "long structures" (ω^2-structures) ([FHV84]), or with the "augmented structures" ([FV85]), which do not need to be of transfinite length in order to represent common knowledge formulas, but which seem rather heavy to handle.

We follow a model-checking approach ([HV91]) as opposed to a deductive approach (as in [Lev84, Lev87, Mor90] for instance): to determine if a given formula is true, an agent evaluates it on his model. This allows us to give a realistic model of interacting agents and of their inference capabilities. For instance our agents have only a partial view of the situation. They are aware of a limited subset of primitive propositions ([Lev84, FH85]) and their current model is a representation of the finite amount of information they have gained thus far.

Université de Liège.

REFERENCES

[Bar88] Jon Barwise. The Situation in Logic-IV: On the Model Theory of Common Knowledge. Research Report CSLI-88-122, CSLI, Mar 1988.

[FH85] R. Fagin and J. Y. Halpern. Belief, Awareness, and Limited Reasoning: Preliminary Report. *In Proceedings of the Ninth IJCAI*, pages 491-501, 1985.

[FHV84] R. Fagin, J. Y. Halpern, and M. Y. Vardi. A Model-Theoretic Analysis of Knowledge: Preliminary Report. Research Report RJ 4373 (47631), IBM Research Laboratory, San Jose, CA 95193, Jul 1984.

[FHV86] R. Fagin, J. Y. Halpern, and M. Y. Vardi. What Can Machines Know? On the Epistemic Properties of Machines. *In AAAI 86 Proceedings*, pages 428-434, Aug 1986.

[FHV91] R. Fagin, J. Y. Halpern, and M. Y. Vardi. A model-theoretic analysis of knowledge. *Journal of the ACM*, 91(2): 382-428, 1991.

[FI86] M. Fischer and N. Immerman. Foundations of Knowledge for Distributed Systems In Halpern [Hal86a], pages 171-186.

[FV85] R. Fagin and M. Vardi. An Internal Semantics for Modal Logic:

Preliminary Report. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, pages 305-315. ACM, 1985.

[Hal86a] J. Y. Halpern, editor. *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning About Knowledge*. Morgan-Kaufmann, 1986.

[Hal86b] J. Y. Halpern. Reasoning About Knowledge: An Overview. In *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 1-18.

[Hin62] J. Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.

[HM84] J. Y. Halpern and Y. Moses. Knowledge and Common Knowledge in a Distributed Environment. In *Proceedings of the 3rd ACM Conference on Principles of Distributed Computing*, pages 50-61, 1984.

[HM85] J. Y. Halpern and Y. Moses. A Guide to the Modal Logics of Knowledge and Belief: Preliminary Draft. In *Proceedings of the Ninth IJCAI*, pages 480-490, 1985.

[HV91] Joseph Y. Halpern and Moshe Y. Vardi. Model Checking vs. Theorem Proving: A Manifesto. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, 1991.

[Kri63] S. Kripke. Semantical Analysis of Modal Logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9: 67-96, 1963.

[Lev84] H. J. Levesque. A Logic of Implicit and Explicit Belief. In *Proceedings of the National Conference on AI*, pages 198-202, Austin, TX, 1984.

[Lev87] H. J. Levesque. All I Know: An Abridged Report. In *Proceedings of the National Conference on AI*, pages 426-431, Seattle, WA, 1987.

[Mor90] L. Morgenstern. A Formal Theory of Multiple Agent Nonmonotonic Reasoning. In *AAAI 90 Proceedings*, pages 538-544, 1990.

[Mos86] Y. Moses. *Knowledge in a Distributed Environment*. Ph. D. dissertation, Stanford University, Stanford, CA 94305, Mar 1986.

[RK86] S. Rosenschein and L. Kaelbling. The Synthesis of Digital Machines With Provable Epistemic Properties. In Halpern [Hal86a], pages 83-98.