Chapter 1

Modeling Agent Epistemic States: An Informal Overview

Recent advances in intelligent agent research (AGENTS, 1997-2001; AAMAS, 2002-2006) have culminated in various agent-based applications that autonomously perform a range of tasks on behalf of human operators. Just to name a few, the kinds of tasks these applications perform include information filtering and retrieval, situation assessment and decision making, and interface personalization. Each of these tasks requires some form of human-like intelligence that must be simulated and embedded within the implemented agent-based application. The concept of *epistemic states* is often used to represent an actual or a possible cognitive state that drives the human-like behavior of an agent. Established traditional artificial intelligence (AI) research in the areas of knowledge representation and inferencing has been transitioned to represent and intelligently reason about the various mental constructs of an agent, including beliefs, desires, goals, intentions, and knowledge (Cohen and Levesque, 1990; Rao and Georgeff, 1991), simulating its human-like cognitive states.

1.1 Models of Agent Epistemic States

The most commonly used models of agent epistemic states are the propositional, probabilistic, and possible world models (Gärdenfors, 1988):

1. An epistemic state in a *propositional model* is represented as a set of propositions that the agent accepts in the epistemic state. These propositions are expressed by sentences in an object language.
2. An epistemic state in a *probabilistic model* is represented by a probability measure defined over the states of some collection of random variables. This probability measure provides the agent's degree of belief about each state of each random variable in the epistemic state.

3. An epistemic state in a *possible world model* is represented by a set of possible worlds that includes the agent's actual state or world, along with any worlds compatible with the agent's knowledge and beliefs. Each world consists of those propositions that the agent accepts in its epistemological worldview.

Irrespective of the model chosen to represent an agent's epistemic state, the state model must always be built via a substantial knowledge engineering effort; this effort may occasionally be aided by machine learning techniques (Mitchell, 1997) to automatically extract knowledge from data. Part of the knowledge engineering effort is the knowledge acquisition process that a scientist or engineer goes through when extracting knowledge from a domain or subject matter expert (such as a medical doctor), to be used by the agent for problem solving in a particular domain (for example, diagnosing diseases). The knowledge gained from this process can then be represented and stored in a computer, using the syntax of logic, frames, or graphical representations, such as semantic networks and probabilistic belief networks. The focus in this text is on logical syntaxes and graphical belief networks (and their amalgamations) that encode an agent's epistemic state that represents knowledge about an uncertain environment. The agent then uses this encoded knowledge for making decisions under uncertainty, that is, to choose between two or more different options. For example, the agent may need to choose to believe in one possibility among many (for example, a patient has cancer or an ulcer), and thus an agent is revising its own belief, or may need to adopt a plan (or take some action), so that the agent is choosing among many possible options (e.g. surgery, chemotherapy, medication) for action.

A note of caution: When we refer to an agent *making a decision*, we do not really mean that agents *always* make decisions autonomously. In some situations, when not enough evidence is available to choose one option over the others, or in cases when some dilemma occurs, agents may very well present the human decision maker with the viable options along with their accumulated supports, backed by explanations. In these situations, agents autonomously provide "decision aids" to human decision makers. Regardless of their role as a decision maker or aid provider, labeling them with the term "agent" is always justified at least due to their autonomous role of routinely searching through huge volumes of data (e.g. Internet web pages) for relevant evidence (a task for which human performance is usually poor) to be used for the decision making process.

In the process of introducing representation and reasoning within various epistemic models, we will often rely on the example of a game (for example, a sporting event) which is scheduled to occur sometime on the current day. In our example, the decision that an agent will usually be trying to make is to determine the status of the game, (that is, whether the game is "on," "cancelled," or "delayed") while preparing to go to town based on accumulated evidence.

1.2 Propositional Epistemic Model

The language in the propositional epistemic model of an agent is usually governed by classical propositional logic; it is uniquely identified by its particular syntax, set of axioms, and inference rules. In general, logic is a systematic study of valid arguments. Each argument consists of certain propositions, called *premises*, from which another proposition, called the *conclusion*, follows. Consider the following argument by an agent to determine the status of a game:

If the field is wet or there is no transportation,

then the game is cancelled. (Premise)

The field is wet. (Premise)

Therefore, the game is cancelled. (Conclusion)

In the above argument, the first two statements are premises. The first is a *conditional* statement, and the second is *assertional*. The third statement is the conclusion, or *argument*. (The term "therefore" is a sign of argument.) The above argument is valid. (Technically, a valid argument is one in which the conclusion must be true whenever the premises are true.)

Valid arguments are studied independently of the premises from which arguments are drawn. This is achieved by expressing valid arguments in their *logical* or *symbolized form*. The valid argument above is symbolized at an abstract level as:

$$P \vee Q \rightarrow R, P \vDash R$$

where the symbols P, Q, and R stand for the proposition "the field is wet," "no transportation," and "the game is cancelled," respectively. So the first premise of the argument, which is a conditional premise, is symbolized as $P \vee Q \rightarrow R$, where '$\rightarrow$' reads "implies". The second premise is symbolized by just the proposition P. The symbol $\vDash$ is the consequence relationship and R is a consequence of the premises. R is arrived at by the use of another argument

$P \vDash P \vee Q$, and applying an *inference rule* of the logic, called *modus ponens*, on $P \vee Q \rightarrow R$ and $P \vee Q$. The inference rule states that given $X \rightarrow Y$ and X is true then it follows that Y is true, where X and Y are arbitrary sentences of the logic.

Although the above argument is valid, it is not known if the outfield is wet in reality; it has just been assumed for the purpose of argument. The validity of an argument is neither concerned with the actual subject matter nor with the truth or falsehood of the premises and conclusion in reality. Mathematical logic does not study the truth or falsehood of the particular statements in the premises and conclusion, but rather focuses on the process of reasoning, i.e. whether or not the assumed truth of the premises implies the truth of the conclusion via permissible symbol manipulation within the context. This phenomenon, to some extent, stimulates interesting philosophical argument on whether an agent, whose epistemic state is modeled using mathematical logic, is really capable of human-like thinking or is just "blindly" manipulating symbols (Searle, 1984).

Often arguments are advanced without stating all the premises, as is evident from the following argument constructed out of the elements of the previous argument:

The field is wet. (Premise)

Therefore, the game is cancelled. (Conclusion)

Due to the lack of an appropriate premise, this example is not a valid argument in the context of classical propositional logic unless the agent intelligently assumes that "If the field is wet then the game is cancelled" is by default a premise. Consider the following example, which has an incomplete premise:

If the field is wet, then the game is cancelled. (Premise)

The field is wet, or the field sensing equipment is inaccurate. (Premise)

Therefore, the game is cancelled. (Conclusion)

The premise of the above argument does not definitely imply that the game is cancelled, since the field may be dry and the field sensing equipment incorrectly signaled a wet field. Hence it is an invalid argument.

There are various kinds of arguments that cannot be stated by propositional logic without an appropriate combinatorial enumeration over the underlying objects. Consider the following example of a valid argument that requires the syntax of the so-called *first-order* logic:

Every wet field is unsuitable for playing a game.

The field at Eden Garden is wet.

Therefore, Eden Garden is unsuitable for playing a game.

Clearly, the assertional second premise does not occur in the conditional first premise, so that propositional logic fails to produce the desired conclusion. One solution to this shortcoming is to produce a specific instantiation of the first premise for the specific Eden Garden field (a sporting venue) and then apply Modus Ponens. But first the conditional premise of the above argument is symbolized as a general sentence:

$$\forall x(Field(x,Wet) \to Unsuitable(x))$$

where x is a variable or placeholder for terms representing things such as field names, *Wet* is a constant symbol representing one possible type of field condition, *Field* is a binary relation or predicate symbol that relates a field to its condition, and *Unsuitable* is a unary relation representing a property of a field. The symbols $\forall$ and $\to$ are read as "for every" (or "for all") and "implies" respectively. The expression $Field(x,Wet)$, which appears on the left , is treated as the *antecedent* of the premise. The expression following the symbol is the *consequent*. Note that there are other ways that the argument can be symbolized depending on the context and inference need. For example, one can define a unary predicate or property *Wet Field* of a field replacing the binary predicate *Field*. (This of course blocks any sort of inference that takes into account the dryness property of the field.)

The second premise of the argument, which is an assertional premise, is symbolized based on the predicate symbol *Field* as follows

$$Field(Eden\,Garden,Wet)$$

A first-order symbolization of the argument now looks like this:

$$\forall x(Field(x,Wet) \to Unsuitable(x)), Field(Eden\,Garden,Wet) \vdash$$
$$Unsuitable(Eden\,Garden)$$

This symbolization is within the framework of *first-order logic*. An axiomatic deduction or inferencing of the conclusion of the argument appears as follows:

Step 1: $\forall x(Field(x,Wet) \rightarrow Unsuitable(x))$ Given Premise

Step 2: $Field(EdenGarden,Wet)$ Given Premise

Step 3: $Field(EdenGarden,Wet) \rightarrow Unsuitable(EdenGarden)$
Axiom on Step 1

Step 4: $Unsuitable(EdenGarden)$ Modus Pones on Steps 2 & 3

The premises in steps 1 and 2 are considered *proper axioms*. The above deduction is a *proof* of the conclusion $Unsuitable(EdenGarden)$ and therefore is a *theorem* that follows from the first-order system with the two proper axioms. Step 3 is derived from step 1 by particularization of the first premise. If the first premise is true for all fields, then it is also true for the field Eden Garden. Step 4 is arrived at by the application of Modus Ponens. These steps demonstrate the basic *axiomatic* theorem-proving approach.

In contrast to the axiomatic theorem proving approach, *resolution* theorem proving is based on the refutation principle. This principle states that if a set of premises is inconsistent then there must be a refutation. Based on this principle, the negation of what is to be proved via a query is assumed as a premise and a refutation is established if the query is to be derived (this very general technique is formally known as *reduction ad absurdum*, or proof by contradiction) . A query could be as general as "Is there any unsuitable playing field?", or could be more specific as in the case of our example, i.e. "Is Eden Garden unsuitable?" The negation of the conclusion in our example arguments context is "It is not the case that Eden Garden is unsuitable" is considered as a premise. The symbolized form of this premise is:

$$\neg Unsuitable(EdenGarden)$$

where $\neg$ is read as "It is not the case." As this assumption is inconsistent with the other two premises, a refutation will result. To do this, first the symbolized premises are converted to their equivalent *clausal form*. In this example the assertional premises are already in clausal form, while the conditional premise can be converted to clausal form to give:

$$\forall x(\neg Field(x,Wet) \vee Unsuitable(x))$$

where $\vee$ is read as "or." By convention, a sentence in clausal form is normally written by removing $\forall x$-type signs from its front and these are assumed by default. In clausal form, the first premise of the argument can now be read as:

Every field is either not wet or is unsuitable for a game which has the same meaning (semantically) as the original form of the first premise. Using this equivalent clausal form, the steps of the refutation of the argument are as follows:

Step 1: $\neg Field(x, Wet) \vee Unsuitable(x)$ Premise

Step 2: $Field(Eden\,Garden, Wet)$ Premise

Step 3: $\neg Unsuitable(Eden\,Garden)$ Negation of query – assumed as premise

Step 4: $Unsuitable(Eden\,Garden)$ Resolution of clause in Steps 1 and 2

Step 5: □ (empty clause) Resolution of clause in Steps 3 and 4

The complement of the atomic premise $Field(Eden\,Garden, Wet)$ in step 2 and the subexpression $\neg Field(x, Wet)$ of the premise in step 1 can be made equal by a *most general unification* $\{x / Eden\,Garden\}$, that is, by substituting *Eden Garden* in place of *x*. The clause $Unsuitable(Eden\,Garden)$ is then obtained by applying the *resolution principle* (Robinson, 1965) to steps 1 and 2. This means applying most general unifiers to both expressions and then merging and canceling the complementary expressions. To complete above derivation, the resolution principle can be applied again to steps 3 and 4. After canceling the complementary atomic expressions, the resultant expression is empty and therefore a refutation of the argument has been arrived at (□ denotes a refutation). This shows that the assumption in step 3 is wrong and therefore its complement is true, that is, Eden Garden is unsuitable for a game.

If-then types of sentences are clauses of a special type that many believe are natural representations of human knowledge. So-called production rules in expert systems are based on the if-then syntax. Many subject matter experts are comfortable expressing their knowledge in this form. In the logic programming community, a set of these rules constitute a program that is guaranteed to be consistent, and most logic programming systems can directly handle such programs without needing any meta-interpreters. Therefore, the use of these rules in building an agent's knowledge is both practical and consistent with established theories.

In building an agent's propositional epistemic state, *logical omniscience* is assumed which means that the agent knows all tautologies and that its knowledge is closed under Modus Ponens. However, omniscience is not realizable in practice since real agents are resource-bounded. Attempts to define knowledge in the presence of bounds include restricting what an agent knows to a set of formulae which is not closed under inference or under all instances of a given axiom.

1.3 Probabilistic Epistemic Model

The knowledge that constitutes an epistemic state of an intelligent software agent will often be imprecise, vague, and uncertain. One major drawback of the propositional approach to modeling epistemic state is that it does not even provide the granularity needed to represent the uncertainty in sentences that we use in everyday life. The truth value of a logical sentence in the propositional approach is either true or false, and there is nothing in between. Therefore no special treatment can be given to a sentence like "If the field is wet then the game is 90% likely to be cancelled." The probabilistic model provides this granularity with probabilities representing uncertainty in sentences.

A *probabilistic model* of epistemic states is represented by a probability measure defined over some space of events represented as random variables. This probability measure provides the agent's degree of belief in each possible world state. A *random variable* is a real-valued function defined over a sample space (that is, the domain of the variable) and its value is determined by the outcome of an experiment, known as an *event*. A *discrete random variable* is a random variable whose domain is finite or denumerable. Probabilistic modeling of an agent's epistemic state involves first identifying the set of random variables in the domain and then developing a joint probability distribution of the identified variables to determine the likelihood of the epistemic state that the agent is in. To illustrate this, consider an example domain to determine the status of a game depending on the field condition and the transport. The random variables in this context are *Game*, *Field*, and *Transport*; they are defined over the sample spaces or states {*on*, *cancelled*, *delayed*}, {*dry*, *wet*}, and {*present*, *absent*} respectively.

The *probability distribution* of a random variable is a function whose domain contains the values that the random variable can assume, and whose range is a set of values associated with the probabilities of the elements of the domain. The *joint probability distribution* of the three discrete random variables *Game*, *Field*, and *Transport* is a function whose domain is the set of triplets (x, y, z), where x, y, and z are possible values for *Game*, *Field*, and *Transport*, respectively, and whose range is the set of probability values corresponding to the ordered pairs in its domain. Therefore, the following probability measure

$$p(Game = cancelled, Field = wet, Transport = present) = 0.8$$

indicates that 0.8 is the probability that the game is cancelled *and* the field is wet *and* transport is present. The probabilities in a joint probability distribution have to be computed from previous data based on the frequency probability approaches. Given a distribution, you can then compute *conditional probabilities* of interest by applying multiplication and marginalization rules. In the frequency approach, probability is derived from observed or imagined frequency distributions. For example, the conditional probability that the game is on given that the field is wet *and* transport is absent is computed as follows:

$$p(Game = on \mid Field = wet, Transport = present) = \frac{p(Game = on, Field = wet, Transport = present)}{\sum\limits_{\substack{x \in \{on, cancelled, \\ delayed\}}} p(Game = x, Field = wet, Transport = present)}$$

In the Bayesian approach, the probability on a particular statement regarding random variable states is not based on any precise computation, but describes an individual's personal judgment (degree of belief) about how likely a particular event is to occur based on experience. The Bayesian approach is more general and expected to provide better results in practice than frequency probabilities alone because it incorporates subjective probability. The Bayesian approach can therefore obtain different probabilities for any particular statement by incorporating prior information from experts. The Bayesian approach is especially useful in situations where no historic data exists to compute prior probabilities. Bayes' rule allows one to manipulate conditional probabilities as follows:

$$p(A \mid B) = \frac{p(B \mid A)\, p(A)}{p(B)}$$

This rule estimates the probability of A in light of the observation B. Even if B did not happen, one can estimate the probability $p(A \mid B)$ that takes into account the prior probability $p(A)$.

Computing the joint probability distribution for a domain with even a small set of variables can be a daunting task. Our example domain with the variables *Game*, *Field*, and *Transport*, and state space sizes 3, 2, and 2 respectively, requires $3 \times 2 \times 2 - 1$ probabilities for the joint distribution. A domain with just 10 binary variables will require $2^{10} - 1$ probabilities. This problem can be mitigated by building *Bayesian belief networks* based on qualitative information in the domain, such as the following:

The field condition and the transport situation together determine the status of the game.

This is simply qualitative information about some dependency between the variables *Field*, *Transport*, and *Game*. We can add even more qualitative information relating the game status and radio commentary as follows:

There is radio commentary when the game is on.

Note that we have no explicit information stated on any dependency between the newly introduced variable *Commentary* and the two variables *Field* and *Transport*.

A Bayesian belief network is a graphical, probabilistic knowledge representation of a collection of random variables (e.g. *Game*, *Field*, *Transport*, and *Commentary*) describing some domain. Nodes of a belief network denote random variables; links are added between the nodes to denote causal relationships between the variables. The topology encodes the qualitative knowledge about the domain. (These descriptions are usually in the form of "causal relationships" among variables.) Conditional probability tables (CPTs) encode the quantitative details (strengths) of the causal relationships, such as the following:

If the field is dry and transport is present there is a 90% chance that the game is on.

The CPT of a variable without any parents consists of just its *a priori* (or prior) probabilities. The belief network of Figure 1-1 encodes the relationships over the domain consisting of the variables, *Game*, *Field*, *Transport*, and *Commentary*; its topology captures the commonsense about the variable dependencies discussed above. Each variable has a mutually exclusive and exhaustive set of possible *states*. For example, *dry* and *wet* are the states of the variable *Field*.

As shown in the figure, the CPT specifies the probability of each possible value of the child variable conditioned on each possible combination of parent variable values. The probability that the game is on given the field is dry and the transport is present is 0.90, whereas the probability that the game is cancelled given the same two conditions is 0.09. Similarly, the probability of radio commentary given the game is on is 0.95. The prior probability that the field is dry is 0.8. Each CPT column must sum up to 1.0 as the states of a variable are mutually exclusive and exhaustive. Therefore, the probability that the game is delayed given that the field is dry and transport is present is $1.0 - 0.9 - 0.09$, or 0.01.

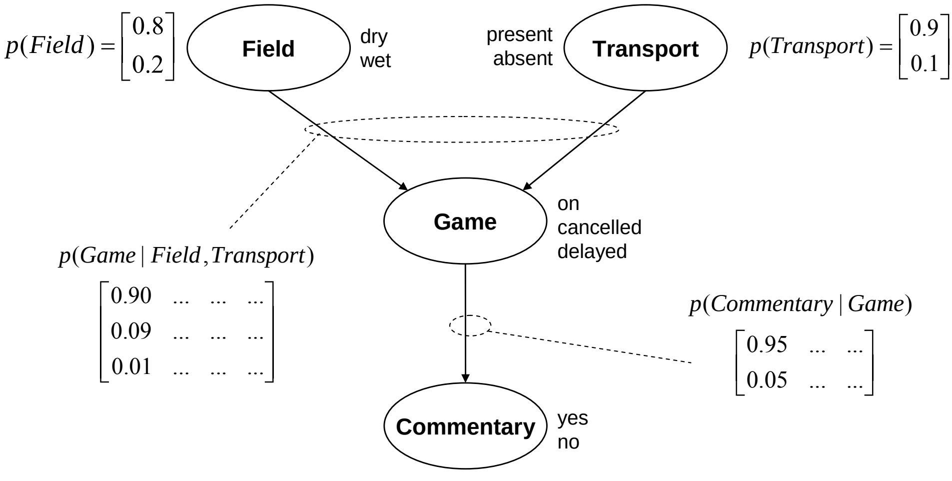

Figure 1-1: Example Bayesian Belief Network

As mentioned above, the structure of a belief network encodes other information as well. Specifically, the lack of links between certain variables represents a lack of direct causal influence, that is, they indicate conditional independence relations. This belief network encodes many independence relations, for example:

$$Commentary \perp \{Field, Transport\} \mid Game$$

which states that *Commentary* is independent of *Field* and *Transport* given *Game*. In other words, once the value of *Game* is known, the values of *Field* and *Transport* add no further information about *Commentary*. Therefore, independence between two nodes is represented by the absence or blocking of links between the two nodes. Whether any dependence between a pair of nodes exists or not is determined by a property called *d-separation*. Conditional independence among variables in a belief network allows factorization of a joint probability distribution, thus eliminating the need for acquiring all the probabilities in a distribution. For example, the joint distribution of the variables in the belief network in Figure 1-1 is factorized as follows:

$$p(Field, Transport, Game, Commentary) =$$
$$p(Field)\, p(Transport)\, p(Game \mid Field, Transport)\, p(Commentary \mid Game)$$

Only 13 (=$1+1+8+3$) probabilities are required to compute the joint probability distribution in this case, as opposed to 23 ($=2 \times 2 \times 3 \times 2 - 1$) when no causal relationship information is given.

When new evidence is posted to a variable (i.e. its state is determined) in a belief network, that variable updates its own state probabilities that constitute its *belief vector* and then sends out messages indicating updated predictive and diagnostic support vectors to its children and parent nodes respectively. The messages are used by the other nodes, which update their own belief vectors, and also propagate their own updated support vectors. The probability distribution of a variable in a belief network at any stage is the *posterior probability* of that variable given all the evidence posted so far. In the case of polytrees, the separation of evidence yields a propagation algorithm (Pearl, 1988) in which update messages need only be passed in one direction between any two nodes after the posting of evidence. See (Jensen, 1996) for a propagation algorithm for the more general case of directed acyclic graphs (DAGs).

It is fairly straightforward to model a decision-making problem in terms of a belief network, where a set of mutually exclusive decision options can be encoded as the states of a random variable representing the decision node within the network. Other attributes from the problem domain are also encoded as random variables connected to the decision node as per their causal relationships. Moreover, belief networks are suitable for synthesizing knowledge at higher levels of abstraction by performing aggregations on low-level evidence. Our experience with various subject matter experts suggests that their thought process during information aggregation is very similar to what the graphical representation of belief networks depicts via their dependencies among random variables.

A belief network is built on the concept of random variables, but a decision making process often involves reasoning explicitly with actions and utilities. *Influence diagrams* are belief networks augmented with decision variables and a utility function to solve decision problems. There are three types of nodes in an influence diagram: *chance nodes* (i.e. belief network nodes), *decision nodes*, and value or *utility nodes*. Using an influence diagram, agents will therefore be able to directly reason about, for example, what loss might be incurred if a decision for proceeding with the game under rainy conditions is made.

1.4 Possible World Epistemic Model

If an agent receives information from a dynamic and uncertain environment, then the agent's understanding about the environment will be based on its own "beliefs," and potential realizations of each of the decision making options will naturally yield a possible world that the agent can transition to. Reasoning within an uncertain environment therefore requires an appropriate underlying knowledge representation language incorporating the concepts of belief and possibility in its semantics, along with a sound inference mechanism. Logical formalisms are especially appealing from the knowledge representation point of view. However, fragments of standard first-order logic in the form of Horn clauses used in most of these formalisms are often inadequate. Researchers have therefore moved to higher-order and non-classical logics. The notion of possibility in classical modal logics and the notion of belief in their epistemic interpretations together provide a useful means for representing an agent's uncertain knowledge. The possibility concept naturally blends with the concept of decision options triggering agents to consider different possible worlds corresponding to those options. The language in the possible world epistemic model is governed by epistemic interpretations of the language of modal logics.

The original purpose behind the revival of modern modal logic by Lewis in his book *A Survey of Symbolic Logic* (Lewis, 1918) was to address issues related to *material implication* of Principia Mathematica (Whitehead and Russell, 1925−1927), by developing the concept of *strict implication*. The language of modal propositional logics extends the language of classical propositional logic with two modal operators $\Box$ (necessary) and $\Diamond$ (possibility). For example, when the necessary operator is applied to the sentence "If the field is wet then the game is cancelled" (symbolized as $\Box(P \to Q)$, where P stand for "the field is wet" and Q stands for "the game is cancelled"), the resultant sentence is read as "It is necessary that if the field is wet then the game is cancelled." The interpretation of $\Box(P \to Q)$ in terms of possible worlds is that in every world that the agent considers possible, if the field is wet then the game is cancelled. Similarly, the interpretation of $\Diamond(P \to Q)$ is that in *at least one* of worlds that the agent considers possible, if the field is wet then the game is cancelled.

Modal epistemic logics (Hintikka, 1962) are instances of modal logics constructed by interpreting necessity and possibility in a suitable manner for handling mentalistic constructs for knowledge and belief. If an agent's epistemic state contains $\langle bel \rangle (P \to Q)$, where the $\langle bel \rangle$ represents the modal operator for belief analogous to the operator $\Box$, then in every world that the agent considers possible (or *accessible* from the current world), the agent believes that if the field is wet, then the game is cancelled. It may so happen that in one of the worlds that the agent considers possible the game is played despite the field being wet. In this case, the agent does not believe in $P \to Q$.

Various modal systems are built by adding axioms to the base modal system $\mathcal{K}$. For example, the system $\mathcal{T}$ is obtained from $\mathcal{K}$ by adding the axiom $\Box F \to F$, which states that whatever is necessary is also true in the current world. Various other modal systems, including the well-known $\mathcal{S}4$, $\mathcal{B}$, and $\mathcal{S}5$ systems are also obtained by adding appropriate axioms to $\mathcal{K}$. The possible world interpretation of the modal operators as explained above is exactly what Kripke (1963) proposed as semantics for various modal systems. The semantics of each modal system imposes some restrictions on the accessibility among possible worlds. For example, the accessibility relation is reflexive in the case of system $\mathcal{T}$. Axiomatic deductions in modal propositional logics are carried out in a manner similar to standard propositional logic, i.e. using a set of inference rules and axioms.

The syntax of modal first-order logics can be obtained simply from the classical first-order logic by the introduction of modalities into the language. Then we can produce first-order modal logics analogous to various propositional modal systems such as $\mathcal{K}$, $\mathcal{D}$, $\mathcal{T}$, $\mathcal{S}4$, $\mathcal{B}$, and $\mathcal{S}5$. In principle, the semantics of these first-order modal systems could simply be obtained by extending the possible world semantics for modal propositional systems. However, lifting up modal propositional systems to their counterparts in modal first-order logics raises some interesting issues regarding the existence of objects in various possible worlds and interactions between the modal operators and quantifiers. For example, in a possible world there may be more or fewer objects than in the current world. Adding the Barcan Formula (BF) (Barcan, 1946) can usually axiomatize first-order modal logics that are adequate for applications with a fixed or non-growing domain:

$$\text{BF: } \forall x \Box F[x] \to \Box \forall x F[x]$$

where $F[x]$ means F has zero or more free occurrences of the variable x. The above axiom means that if everything that exists necessarily possesses a certain property F then it is necessarily the case that everything possesses F. But in the cumulative domain case there might be some additional objects in accessible worlds that may not possess the property F.

In general, there is no equivalent to the standard, first-order logic, clausal representation of an arbitrary modal formula where each component in the clause is one of the atomic modal formulae P, $\neg P$, $\Diamond P$, and $\Box P$. But a modal formula can be transformed into an equivalent formula in first-order like world-path syntax, which can then be transformed into a set of clauses in the same way as in the first-order logic. For example, the formula $\Diamond \forall x P(x)$ is transformed into

$\forall x P([0a], x)$, meaning that from the initial world 0 there is a world *a* accessible from 0 such that for all *x* $P(x)$ holds, where *P* is interpreted in the world *a*. On the other hand, $\Box\forall x P(x)$ is transformed into $\forall u \forall x P([0u], x)$, meaning that from the initial world 0 and for every world *w* that is accessible from *a* for all *x*, $P(x)$ holds, where *P* is interpreted in the world *u*. Therefore, the symbolization of "It is necessary that every wet field is unsuitable for playing game":

$$\Box\forall x(Field(x, Wet) \rightarrow Unsuitable(x))$$

can be transformed into the following clause in world-path syntax:

$$\forall u \forall x(\neg Field([0u], x, Wet) \vee Unsuitable([0u], x))$$

Resolution theorem proving in this setting (Ohlbach, 1988) is carried out on a set of clauses of the above form in a way similar to standard first-order resolution theorem proving except for the unification of terms like $[0u]$, which is carried out by taking into account the type of accessibility relation among possible worlds.

The notion of possibility in classical modal logics is useful for representing and reasoning about the uncertain knowledge of an agent. However, this coarse grain representation of the knowledge of possibilities about assertions through this logic is quite inadequate for practical applications. That is to say, if two assertions are possible in the current world, they are indistinguishable in the modal formalism even if an agent knows that one of them is true in twice as many possible worlds as compared to the other one. Therefore, any epistemic interpretation of the modal operators for necessity and possibility would fail to incorporate the notion of an agent's "degree of belief" in something into its epistemic states.

Therefore, we developed an extended formalism of modal epistemic logic to allow an agent to represent its *degrees of support* about an assertion. The degrees are drawn from qualitative and quantitative dictionaries that are accumulated from the agent's *a priori* knowledge about the application domain. We extend the syntax of traditional modal epistemic logic to include an indexed modal operator $\langle sup_d \rangle$ to represent an agent's degrees of support about an assertion. In this way, the proposed Logic of Agents Beliefs ($\mathcal{LAB}$) can model an argument that merely supports an assertion (to some extent *d*), but does not necessarily warrant an agent committing to believe in that assertion. The extended modal logic is given a modified form of possible world semantics by introducing the concept of an accessibility *hyperelation*.

1.5 Comparisons of Models

Table 1-1 presents a comparison of various features of the three epistemic models presented in the last three sections. Detailed pros and cons of individual decision making frameworks under these models (e.g. logical rules and belief networks) and their suitability for different application types are discussed in their corresponding chapters.

	Underlying formalism	**Qualitative information**	**Uncertainty handling**	**Inferencing mechanism**	**Efficiency enhancement**
Propositional	Classical logics	Logical sentences	None	Theorem proving	Horn clauses, resolution
Probabilistic	Probability theory	Causal networks	Fine-grained	Evidence propagation	Conditional independence
Possible worlds	Modal logics	Modal sentences	Coarse-grained	Theorem proving	Modal resolution

Table 1-1: Comparison of features of epistemic models

The underlying formalism for the propositional epistemic model is the classical propositional and first-order logics, whereas various systems of modal logics formalize the possible world epistemic model. The probabilistic epistemic model, as its name suggests, is formalized by the theory of probability, especially Bayesian probability.

Qualitative domain information in the propositional epistemic model is expressed by user-defined relations among domain objects and by logical connectives among various assertions, producing sentences. Qualitative domain information in the probabilistic epistemic model is in the form of causal relationships among various domain concepts represented as random variables, producing belief networks. Qualitative domain information in the possible world epistemic model is also expressed by user-defined relations and logical connectives. Additionally, modalities are allowed to apply to sentences to produce qualitative domain information.

The truth-value of an assertion in the classical logics is Boolean (either true or false), and therefore the propositional epistemic model is unable to handle uncertainties in assertions. Modal logics provide a very coarse-grained uncertainty representation through their modal operator for possibility. An agent's belief in an assertion does not imply that the assertion is true in the world the agent is in. The belief rather describes the agent's uncertain assumption of the true nature of the assertion. Uncertainty handling in the probabilistic epistemic model is fine-grained in the sense that an agent can choose any number from the dense real-valued dictionary [0,1] of probability as its degree of belief in a random variable state.

The underlying inferencing mechanism for deriving implicit assumptions in the propositional and possible world epistemic models is based on axiomatic theorem proving. The inferencing mechanism for deriving prior and posterior probability distributions in the probabilistic epistemic model is based on the application of Bayes' rule, or takes the form of evidence propagation when models are available as causal belief networks.

To enhance inferencing efficiency in the propositional epistemic model, the resolution theorem proving technique is employed. The possible world epistemic model employs the same concept in the context of modal logics. To further enhance inferencing efficiency, sentences in the propositional epistemic model are often expressed in the Horn clause syntax, which is less general than full first-order language but expressive enough to deal with most common applications. Expert system rules are just positive Horn clauses. Various conditional independences among random variables are assumed in the probabilistic epistemic model.

1.6 *P3* Model for Decision-Making Agents

So far we have described three models of epistemic states, namely propositional, possible world, and probabilistic, and compared their features. Propositional and possible world models are ideal for representing agent knowledge at a higher level natural language-like syntax and for thinking semantics in terms of possible worlds, and the probabilistic model provides the required granularity with probabilities representing uncertainty in sentences that we use in everyday life, but not one of these models by itself can sufficiently capture an agent's epistemic states. Consequently, there is a need for suitably combining these three models into one. The focus of the final part of this book is to develop such an integrated model of agent epistemic states, called *P3* (Propositional, Probabilistic, and Possible World). The *P3* model is based on a logical language with embedded modalities and probabilities and is particularly suitable for making decisions in the context of choosing between possible courses of action or between possible assertions for belief revision. An intelligent agent continuously performs this kind of cognitive task when making decisions.

The human decision-making process can be regarded as a complex information processing activity. According to (Rasmussen, 1983), the process is divided into three broad categories that correspond to activities at three different levels of complexity. At the lowest level is skill-based sensorimotor behavior, representing the most automated, largely unconscious level of skill-based performance such as deciding to brake upon suddenly seeing a car ahead. At the next level is rule-based behavior exemplified by simple procedural skills for well-practiced, simple tasks, such as inferring the condition of a game-playing field based on the current rainy weather. Knowledge-based behavior represents the most complex cognitive processing. It is used to solve difficult and sometimes unfamiliar problems, and for making decisions that require dealing with various factors and uncertain data. Examples of this type of processing include determining the status of a game given the observation of transport disruption.

The proposed $\mathcal{P}3$ model, grounded in the logic $\mathcal{LAB}$ introduced earlier, supports embedding the human decision-making process within an agent at the knowledge base level by providing suggestions on alternative courses of action, and helping to determine the most suitable one. Human decision-makers often weigh the available alternatives and select the most promising one based on the associated pros and cons. The $\mathcal{P}3$ model will represent these pros and cons as logical sentences with embedded probabilities as follows:

$$\langle bel \rangle Heavy\,Rain \rightarrow \langle sup_{0.7} \rangle Cancelled$$

The above sentence can be interpreted as follows: if the agent believes that it rained heavily then it asserts that there is a 70% chance (equivalently, generates an amount of support 0.7) that the game will be cancelled. An agent may obtain evidence from different sources as support for the cancellation, as well as support against the cancellation, as follows:

$$\langle bel \rangle Club\,Financial\,Crisis \rightarrow \langle sup_{0.6} \rangle \neg Cancelled$$

The above sentence states that if the club is in financial crisis then there is a 60% chance that the cancellation will be avoided. These types of $\mathcal{P}3$ sentences that provide support both for and against various decision options, constitute *arguments* used by the agent to solve a decision problem. Such an argumentation-based decision-making framework has been developed in (Das et al. 1997; Fox and Das, 2000).

The set of evidence for and against a certain decision option F must be aggregated to come up with the overall support for F. This aggregation process can be implemented using Dempster's combination rule within the framework of the Dempster-Shafer theory of belief functions (Shafer, 1976; Yager et al., 1994), but this requires supports to be interpreted as a mass distribution, as opposed to a probability distribution, along with an appropriate evidence independence assumption. If supports are drawn from the dictionary of probability, any aggregation process must be consistent with the theory of probability, which is a special case of Dempster-Shafer theory. The belief network technology can perform such an aggregation process on arguments "for" options, but not on arguments "against" an option. Its propagation mechanism is consistent with the theory of probability. To illustrate this, let us explain the use of belief network technology for aggregation by considering the following two sentences supporting the option *Cancelled*:

$$\langle bel \rangle Heavy\, Rain \rightarrow \langle sup_{0.7} \rangle Cancelled$$
$$\langle bel \rangle Terrorist\, Threat \rightarrow \langle sup_{0.9} \rangle Cancelled$$

The agent will now go through an aggregation process to decide whether or not to believe in game cancellation. One approach is based on the assumption that the heavy rain and terrorist threat conditions are a set of independent conditions that are likely to cause the cancellation of game and that this likelihood does not diminish or increase when several of these conditions prevail simultaneously. This is the effect of the *noisy-or* technique in belief networks, which is usually applied to generate conditional probabilities in large tables for nodes with multiple parents. In many applications this approach makes sense, as experts will often enumerate a list of factors causally influencing a particular event along with their associated uncertainty. The belief network in Figure 1-2 is constructed out of the two sentences mentioned above for aggregating all evidence for the cancellation of the game.

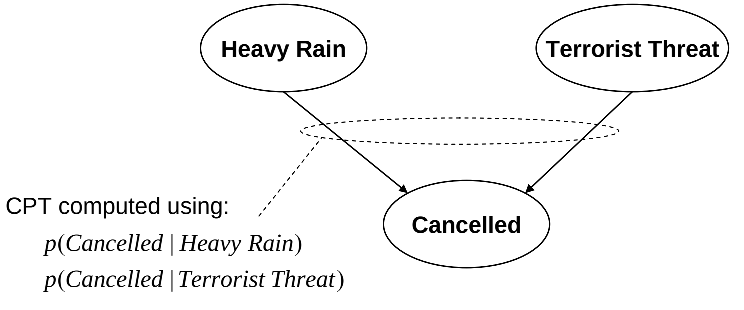

Figure 1-2: Aggregation with belief networks

As shown in the figure, there are two parents of the node *Cancelled*, corresponding to the two sentences. The sentences provide the following probabilities:

$p(C \mid HR) = 0.7$

$p(C \mid TT) = 0.9$

where $C \equiv Cancelled$, $HR \equiv Heavy\,Rain$, $TT \equiv Terrorist\,Threat$

By applying the noisy-or technique, the following probabilities for the table are computed:

$p(C \mid HR, \neg TT) = 0.7$

$p(C \mid HR, TT) = p(C \mid HR) + p(C \mid TT) - p(C \mid HR) \times p(C \mid TT) = 0.97$

$p(C \mid \neg HR, TT) = 0.9$

$p(C \mid \neg HR, \neg TT) = 0$

As evidence is posted into the network, the posterior probability (or the agent's degree of belief in the cancellation node) changes. An agent may decide to accept the cancellation of the game as its belief and $\langle bel \rangle Cancelled$ will be added to its database.

To summarize the *P3* model, we use the syntax of modal propositional logics for representing arguments, and include probabilities to represent their strengths. The use of modal logics allows agents to consider decision options in terms of the intuitive possible world concept, each of which is the result of committing to one decision option. When aggregating a set of arguments to choose a decision option, we can either apply the Dempster-Shafer theory of belief functions with an appropriate evidence independence assumption, or apply the belief network evidence propagation technique on a restricted set of arguments.

Before we delve into the details of various agent epistemic models, it is worth pointing out the role of belief revision in the context of our proposed agent decision-making paradigm. As an agent perceives an environment, it revises its own belief about the environment, which amounts to updating some of its own knowledge representing its own epistemic state. The monotonic nature of logic-based modeling of epistemic states does not automatically allow such a belief revision process. The agent can potentially continue accumulating knowledge as long as the knowledge base is not inconsistent (i.e. deriving a formula F and its negation). A special inference rule or some kind of meta-level reasoning is required for belief revision in case an inconsistency arises between the observation and what is already in the epistemic state.

The belief network-based probabilistic epistemic model is more flexible in accommodating observations from uncertain environments because a consistent probability assignment can always be obtained to yield a consistent epistemic state unless it is a hard inconsistency (for example, the agent already believes in *F* with degree of support 1.0, but subsequently observed its negation with degree of support 1.0) or unless a soft probabilistic consistency concept is incorporated as in (Das and Lawless, 2005). Various propagation algorithms for computing posterior probabilities of random variables based on their priors and observations do these assignments. Our approach to decision making here deals with a "snapshot" of the agent epistemic state rather than its evolution over time. Moreover, the argumentation based $\mathcal{P}3$ model does not require an epistemic state to be consistent because a hard inconsistency will at worst yield a dilemma amongst decision options.