An Architecture for Rational Agents

J.W. Lloyd and T.D. Sears

Computer Sciences Laboratory
Research School of Information Sciences and Engineering
Australian National University
{jwl,timsears}@csl.anu.edu.au

Abstract. This paper is concerned with designing architectures for rational agents. In the proposed architecture, agents have belief bases that are theories in a multi-modal, higher-order logic. Belief bases can be modified by a belief acquisition algorithm that includes both symbolic, on-line learning and conventional knowledge base update as special cases. A method of partitioning the state space of the agent in two different ways leads to a Bayesian network and associated influence diagram for selecting actions. The resulting agent architecture exhibits a tight integration between logic, probability, and learning. This approach to agent architecture is illustrated by a user agent that is able to personalise its behaviour according to the user's interests and preferences.

1 Introduction

Our starting point is the conventional view that an agent is a system that takes percept sequences as input and outputs actions [10, p.33]. Naturally enough, on its own, this barely constrains the agent architectures that would be feasible. Thus we add a further constraint that agents should act rationally, where a rational agent is defined as follows [10, p.36]:

> "For each possible percept sequence, a rational agent should select an action that is expected to maximise its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has".

Various specific rationality principles could be used; we adopt the well-known principle of maximum expected utility [10, p.585] (namely, a rational agent should choose an action that maximises the agent's expected utility). This implies that, for the intended applications, it is possible for the agent designer to specify utilities for all states (or, at least, the agent has some means of acquiring these utilities). The proposed architecture includes a decision-theoretic component involving utilities and Bayesian networks that implements the principle of maximum expected utility.

Another major component of the architecture is a model of the environment (or, at least, a model of enough of the environment in order to be able to effectively select actions). This model has two parts: state and beliefs. The state

records some information that helps the agent mitigate the well-known problems that arise from partial observability. The beliefs are random variables defined on the state; evidence variables assist in the selection of actions and result variables assist in the evaluation of the utility of states. Beliefs are expressed in a multi-modal, higher-order logic.

Implicit in the concept of rationality is the notion that agents should make every effort to acquire whatever information is needed for action selection. This implies that agents should have a learning component that allows them to improve their performance (which includes adapting to changing environments). The proposed architecture incorporates a learning component for exactly this purpose.

The resulting agent architecture exhibits a tight integration between logic, probability, and learning. In the terminology of [10, p.51], agents employing this architecture are model-based, utility-based, learning agents.

2   Agent Architecture

In this section, we describe the agent architecture.

We consider an agent situated in some environment that can receive percepts from the environment and can apply actions to the environment. Thus the agent can be considered to be a function from percepts to actions; an agent architecture provides the definition of this function.

Let S denote the set of states of the agent. Included in a state may be information about the environment or something that is internal to the agent. It is also likely to include the agent's current intention or goal, that is, what it is currently trying to achieve. The state may be updated as a result of receiving a percept. For example, a user agent may change its intention due to a request from the user.

As well as some state, the agent's model includes its belief base, which can also be updated. In the proposed architecture, belief bases have a particular form that we introduce below motivated by the desire to make action selection as effective as possible.

Let A denote the set of actions of the agent. Each action changes the current state to a new state (possibly non-deterministically). The agent selects an action that maximises the expected utility. Executing the action involves applying the action to the environment and/or moving to a new state.

Summarising the description so far, here is a high-level view of the main loop of the agent algorithm.

```
loop forever
    get percept
    update model
    select action
    put action.
```

We now concentrate on the action selection part of this loop. In practical applications, the set of states may be very large; in this case, it may be impractical to try to select the action by explicitly dealing with all these states. An obvious idea is to partition the state space so that the behaviour of an action is somehow consistent over all the states in an equivalence class [1]. We exploit this idea in what follows.

For that we will need some random variables. It will be helpful to be quite precise about their definition. Random variables are functions; this fact and the probability space that they are defined on will be important in the description of the agent architecture. For simplicity and because it is the case of most interest in applications, we confine the discussion to random variables that are discrete.

Definition 1. *A* random variable *is a measurable function $X : \Omega \to \Xi$, where $(\Omega, \mathcal{A}, p)$ is a probability space and $(\Xi, \mathcal{B})$ is a measurable space. The probability measure $X^{-1} \circ p$ on $\mathcal{B}$ is called the* distribution *(or* law*) of X.*

The probability space of interest here is the product space $S \times A \times S$ which is assumed to have some suitable σ-algebra and probability measure $p(\cdot, \cdot, \cdot)$ defined on it. Intuitively, $p(s, a, s')$ is the probability that applying action a will result in a transition from state s to state s'. By conditioning on states and actions, for each state $s \in S$ and action $a \in A$, a transition probability distribution $p(\cdot \mid s, a)$ is obtained.

There are three projections defined on $S \times A \times S$. The first projection *initial* : $S \times A \times S \to S$ is defined by

$$initial(s, a, s') = s,$$

the second projection *action* : $S \times A \times S \to A$ is defined by

$$action(s, a, s') = a,$$

and the third projection *final* : $S \times A \times S \to A$ is defined by

$$final(s, a, s') = s',$$

for each $(s, a, s') \in S \times A \times S$. Each projection induces a probability measure on their respective codomains that makes the codomains into probability spaces.

We will be interested in two different sets of random variables on S. These are the *evidence variables* $E_i : S \to V_i$ $(i = 1, \dots, n)$ and the *result variables* $R_j : S \to W_j$ $(j = 1, \dots, m)$. Two sets of random variables can now be defined on the probability space $S \times A \times S$, as follows. For each evidence variable E_i, consider the random variable $initial \circ E_i : S \times A \times S \to V_i$. Similarly, for each result variable R_j, consider the random variable $final \circ R_j : S \times A \times S \to W_j$. Let X denote $(initial \circ E_1, \dots, initial \circ E_n, action, final \circ R_1, \dots, final \circ R_m)$. All this can now be put together to obtain the random variable

$$X : S \times A \times S \to V_1 \times \cdots \times V_n \times A \times W_1 \times \cdots \times W_m,$$
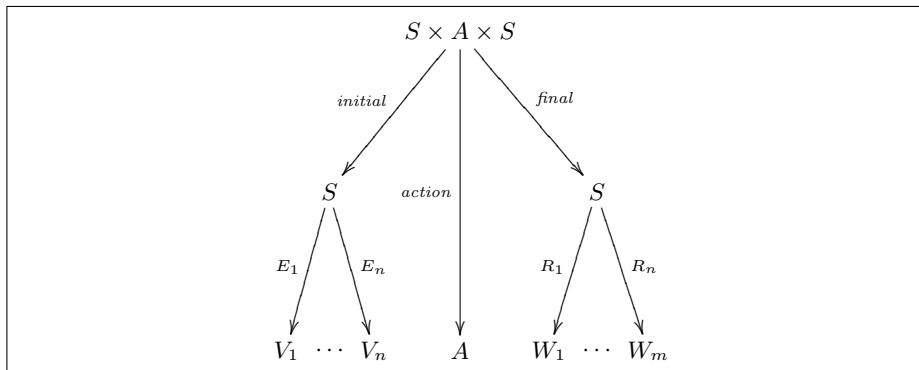
Fig. 1. Evidence and result variables

which is illustrated in Figure 1. The distribution of X is given by $X^{-1} \circ p$, where p is the probability measure on $S \times A \times S$.

In the following, we will refer to an element of $V_1 \times \cdots \times V_n$ as an *evidence tuple* and an element of $W_1 \times \cdots \times W_m$ as a *result tuple*.

The distribution of X is illustrated in the influence diagram in Figure 2. Here the Bayesian network given by the evidence and result variables is extended into an influence diagram by adding the random variable *action* for the action selected and a utility node to indicate that it is the result variables that contribute to the utility. Each node in the Bayesian network has an associated (conditional) probability table (that is not shown in Figure 2). Note that, in general, it is possible for there to be dependencies amongst the *initial* $\circ E_i$ and amongst the *final* $\circ R_j$.

The evidence and result variables play an important role in action selection. One can think of these as features of the states. However, each class of features serves a different purpose. The *evidence variables are chosen so as to assist the selection of a good action*, whereas the *result variables are chosen so as to provide a good evaluation of the resulting state*. It will be convenient to insist that included amongst the evidence variables are boolean random variables that describe the pre-condition (if any) of each action.

These random variables will be functions with definitions in the belief base. We can now be more precise about the definition of belief bases.

Definition 2. *A belief base is a theory consisting of the definitions of the evidence variables and (some of) the result variables.*

The implication of this definition is that the belief base should contain the definitions of the evidence variables, (some of) the result variables, and any subsidiary functions used to define these, and that is *all* it need contain – there is no other use for functions in the belief base. This fact provides strong guidance during the design of the agent. Observations of the result variables will be needed but it will not be essential for the agent to know their (full) definitions. An

Fig. 2. Influence diagram

application for which the definition of a result variable is not known but for which the variable can be adequately observed is given below.

To express belief bases, we employ the logic from [6] which is a multi-modal version of the higher-order logic in [5] (but leaving out polymorphism). We refer the reader to these works for the details, and simply make some motivational remarks here.

The modal aspect of the logic is easily motivated by the need to express epistemic notions, in this case beliefs [4]. We need a *multi*-modal logic because each agent in a multi-agent system has its own modality. Thus a formula of the form $\Box_i \varphi$ is interpreted as 'agent i believes φ'.

An important feature of higher-order logic is that it admits functions that can take other functions as arguments and thus has greater expressive power for knowledge representation than first-order logic. This fact is exploited throughout the architecture, in the use of predicates to represent sets and in the predicate rewrite systems used for learning, for example.

In applications there are typically many individuals that need to be represented. For example, with agent applications, the state is one such individual. With logic as the representation language, (closed) terms represent individuals. In [5], a class of terms, called basic terms, is identified for this purpose. The inductive definition of basic terms comes in three parts. The first part uses data constructors to represent numbers, lists, and so on. The second part uses abstractions to represent sets, multisets, and so on. The third part uses tuples

to represent vectors. The class of basic terms provides a rich class of terms for representing a variety of structured individuals.

As shown in [5], and also [6], a functional logic programming computational model can be used to evaluate functions that have equational definitions in the logic. Furthermore, [6] provides a tableau theorem proving system for the multi-modal, higher-order logic.

In [7], a method of belief acquisition is introduced for belief bases that are theories in the logic. The belief acquisition algorithm includes both symbolic, on-line learning and conventional knowledge base update as special cases. The key idea of the algorithm is to introduce two languages, the training language and the hypothesis language. By carefully controlling these languages it is possible to have belief acquisition that at one extreme directly incorporates new beliefs into the belief bases (as for a conventional knowledge base update algorithm) and at the other extreme is a conventional learning algorithm that learns definitions of functions that generalise to new individuals. The algorithm itself is a (somewhat generalised) decision-list learning algorithm, as introduced in [9]. Thus definitions of functions in belief bases are (logical forms of) decision lists.

We now return to the probabilistic component of the architecture. Over the last 20 years or so there has been a huge effort to find efficient algorithms for doing all kinds of tasks such as marginalising, conditioning, and learning in graphical probabilistic models with large numbers of random variables. (See, for example, [3] and [8].) All of this work is directly relevant to the agent context and we hope to exploit it in future. For the moment, we are primarily interested in making decisions which does not require the full graphical model and the applications of current interest have small graphical models anyway, so we simply make a few remarks below about this context. On the other hand, a new problem is raised by the agent applications in that the definitions of the random variables are generally changing and this leads to some complications, as noted below, that deserve a deeper investigation.

By conditioning on the evidence variables and action variable in Figure 2, the table in Figure 3 is obtained. There is a row in this table for every combination of an evidence tuple together with an action, except that rows for which the precondition (if any) of the action has value $\perp$ are deleted. There is a column under $final \circ (R_1, \ldots, R_m)$ for every result tuple. For each combination of evidence tuple and action, there is an associated transition probability distribution that gives, for each possible result tuple, the probability of reaching that result tuple. The last row records the utilities for each result tuple. To compute the expected utility $EU(e, a)$ of action a applied to evidence tuple e, we proceed as follows. Suppose that e and a appear together in the ith row of the table. Then the expected utility is

$$EU(e, a) = \sum_{j=1}^{l_1 \cdots l_m} p_{i,j} \times u_j.$$

A policy for an agent is a mapping $policy : S \rightarrow A$. The policy is extracted from the conditioned influence diagram in the usual way. Given a state s, the

		$final \circ (R_1, \ldots, R_m)$			
$initial \circ (E_1, \ldots, E_n)$	$action$	$(w_{1,1}, \ldots, w_{m,1})$	$(w_{1,2}, \ldots, w_{m,1})$	$\ldots$	$(w_{1,l_1}, \ldots, w_{m,l_m})$
$(v_{1,1}, \ldots v_{1,n})$	a_1	$p_{1,1}$	$p_{1,2}$	$\ldots$	$p_{1,l_1 \cdots l_m}$
$(v_{2,1}, \ldots v_{2,n})$	a_2	$p_{2,1}$	$p_{2,2}$	$\ldots$	$p_{2,l_1 \cdots l_m}$
$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$
$(v_{k,1}, \ldots v_{k,n})$	a_k	$p_{k,1}$	$p_{k,2}$	$\ldots$	$p_{k,l_1 \cdots l_m}$
		u_1	u_2	$\ldots$	$u_{l_1 \cdots l_m}$

Fig. 3. Influence diagram conditioned on the evidence and action variables

action a selected by *policy* is the one for which $EU((E_1(s), \ldots, E_n(s)), a)$ is a maximum. The case when the state s is not known exactly but is given by a distribution (the 'belief state' case [10]) is handled by the obvious generalisation of this. This policy thus implements the principle of maximum expected utility.

The agent architecture exhibits a tight integration between logic and probability, since the random variables in the network have definitions given by the logic. Furthermore, for agents in dynamic situations, the definitions of evidence variables will need to be modified from time to time. In the first illustration of the next section, we discuss a personalisation application for which there are three evidence variables, one of which is learned and the other two are subject to updating. In effect, the evidence variables are tailored to the interests and preferences of the user and this will be seen to be essential for selecting good actions.

The approach taken here to agent architecture assumes that it is possible to specify the utilities. Generally, assigning utilities to a very large number of states in advance is a difficult task. However, with the right choice of result variables, the task can become much easier; this is illustrated with the application in the next section. The modelling challenge for the agent designer is to find result variables so that all states that map to the same result tuple really do have the same (or very similar) utility. Thus, in essence, the task is to find a good piecewise constant approximation of the utility function. While not every application will succumb to this approach, it seems that there is a sufficiently large subset of agent applications which does and therefore the approach is useful.

Having settled on the actions, and evidence and result variables, and having specified the utilities, the only other information needed to obtain a policy is the collection of transition probability distributions $p(. \mid e, a)$ in Figure 3. How are these obtained? Essentially all that needs to be done is to observe triples of the form (e, a, r), where e is an evidence tuple, a is an action, and r is the corresponding result tuple. The observation of each such triple increments a count kept at the corresponding entry in the table of transition probabilities. More generally, r is not uniquely determined by e and a but is given by a probability distribution. In this case, the increment is proportioned over the result tuples according to this distribution. These counts can then be normalised over each row to obtain a probability distribution.

Typically, the evidence values in an observation (e, a, r) are obtained from some state using the definitions of the evidence variables E_i and the result values are obtained either from some state using the definitions of result variables R_j or from an oracle (the user, for example). Provided these definitions and/or the oracle remain unchanged, the approach to obtaining the transition probability distributions when the observations start from states or involves an oracle works fine. However, in the general case, this assumption does not hold and it will be necessary to do some recalculation of earlier observations. Later we give an application where all the evidence variables change from time to time. Furthermore, one of the result variables represents the user who may very well give different values for the result variable (simply because their interests and preferences have changed). We consider each of these cases in turn.

Consider first the case that the definition of an evidence variable E_i changes. This is handled by the belief acquisition algorithm. The aspect of this algorithm that is relevant to the discussion here is that the definitions of acquired functions are decision lists. When the function definition changes, there is some distinguished decision node in the decision list above which the definition is unchanged and below which it is changed. Each decision node of the decision list contains the states that satisfy all the conditions down to that point in the decision list. Thus when the definition of E_i changes, this distinguished node is identified and all states that occur there have their E_i value recomputed. If the new value is the same as the old, nothing needs to be done. If the new value is different to the old, then the effect is to move (proportions of) an increment from one or more positions in one row of the table of transition probabilities to the corresponding positions in another row. Clearly it is easy to do the same for several E_i's that are changing at the same time.

Consider next the case when the definition of a result variable R_j changes. If the definition of R_j is known to the agent then we can proceed similarly to the preceding paragraph except that (proportions of) an increment are now moved in the same row from some positions to others. The other case is when the value of R_j is obtained from an oracle but its definition is unknown to the agent. In this case, the agent must at least know the relevant state. Then it is only necessary to check whether this state had been considered previously and, if so, whether its R_j value is the same. If it is the same, nothing need be done. If it is different, then an increment has to be moved appropriately along a row from one position to another. It is easy enough to handle the case of several evidence and result variables all changing at the same time.

3   Illustration

We now consider an illustration of the agent architecture. This is concerned with applying machine learning techniques to building user agents that facilitate interaction between a user and the Internet. It concentrates on the topic of personalisation in which the agent adapts its behaviour according to the in-

terests and preferences of the user. There are many practical applications of personalisation that could exploit this technology.

This illustration is set in the context of an infotainment agent, which is a multi-agent system that contains a number of agents with functionalities for recommending movies, TV programs, music and the like, as well as information agents with functionalities for searching for information on the Internet. Here we concentrate on the TV recommender as a typical such agent. More detail (but not the decision-theoretic discussion that follows) is given in [2].

A detailed description of the most pertinent aspects of the design of the TV recommender is now given. The knowledge representation aspects of the TV recommender are presented first; these mainly concern the states, actions, evidence variables, and result variables.

First we introduce the types that will be needed and the data constructors corresponding to these types. We will need several standard types: Ω (the type of the booleans), Nat (the type of natural numbers), Int (the type of integers), and $String$ (the type of strings). Also $List$ denotes the (unary) list type constructor. Thus, if α is a type, then $List\ \alpha$ is the type of lists whose elements have type α.

We introduce the following type synonyms.

$$State = Occurrence \times Status$$
$$Occurrence = Date \times Time \times Channel$$
$$Date = Day \times Month \times Year$$
$$Time = Hour \times Minute$$
$$Title = String$$
$$Subtitle = String$$
$$Duration = Minute$$
$$Synopsis = String$$
$$Program = Title \times Subtitle \times Duration \times (List\ Genre) \times Classification \times Synopsis$$
$$Year = Nat$$
$$Month = Nat$$
$$Day = Nat$$
$$Hour = Nat$$
$$Minute = Nat$$
$$Text = List\ String.$$

The data constructors for the type $Status$ are as follows.

$$Unknown, Yes, No : Status.$$

The meaning of $Unknown$ is that a recommendation (about a program having a particular occurrence) hasn't yet been made, Yes means that it has a positive recommendation, and No means that it has a negative recommendation.

Here are the data constructors for the types *Channel*, *Genre*, and *Classification*.

$ABC, Adventure\_1, Animal\_Planet, Arena, Biography, BBC\_World,$

$\vdots$

$TCM, Tech\_TV, Travel, TV1, UK\_TV, W, World\_Movies : Channel$

$ActionAdventureGroup, Adult, Animals, Animated, Art, ArtsMusicLiving,$

$\vdots$

$War, Watersports, Weather, Western, WesternGroup, Wrestling : Genre$

$Y7, Y, G, MA, M14, M, NA : Classification.$

There are 49 channels, 115 genres and 7 classifications.

There is a type *Action* with data constructors given by

$RecommendYes, RecommendNo : Action.$

The action *RecommendYes* is a positive recommendation, while *RecommendNo* is a negative recommendation. The action *RecommendYes* takes a state $(o, \textit{Unknown})$, where o is some occurrence, and produces the new state $(o, \textit{Yes})$. Similarly, the action *RecommendNo* takes a state $(o, \textit{Unknown})$ and produces the new state $(o, \textit{No})$.

In the following, the definitions of various functions will appear. These are (mainly) in the belief base of the TV agent. To indicate that they are beliefs of the TV agent, the necessity modality $\Box_t$ is used. Thus, if φ is a formula, then the meaning of $\Box_t \varphi$ is that 'φ is a belief of the TV agent'. Other agents in the multi-agent system have their own necessity modality; for example, the modality for the diary agent is $\Box_d$. There is also a base of common beliefs accessible to all the agents for which the necessity modality in $\Box$. Interaction axioms allow one agent to access the beliefs of another agent [6]. So, for example, there are interaction axioms that allow the TV agent to access the beliefs of the diary agent and also the common beliefs.

There are three projections on $State \times Action \times State$. The first projection is

$initial : State \times Action \times State \rightarrow State$

$\Box_t \forall_{State} s. \forall_{Action} a. \forall_{State} s'.$

$\quad ((initial\ (s, a, s')) =_{State} s).$

The second projection is

$action : State \times Action \times State \rightarrow Action$

$\Box_t \forall_{State} s. \forall_{Action} a. \forall_{State} s'.$

$\quad ((action\ (s, a, s')) =_{Action} a).$

The third projection is

$$final : State \times Action \times State \rightarrow State$$

$$\Box_t \,\forall_{State} s. \,\forall_{Action} a. \,\forall_{State} s'.$$
$$((final \,(s, a, s')) =_{State} s').$$

Now we turn to the evidence variables. To define these, a number of subsidiary functions are needed and so we start there.

The agent has access via the Internet to a TV guide (for the next week or so) for all channels. This database is represented by a function tv_guide having signature

$$tv\_guide : Occurrence \rightarrow Program.$$

Here the date, time and channel information uniquely identifies the program and the value of the function is (information about) the program itself. The TV guide consists of (thousands of) facts like the following one.

$$\Box_t \,((tv\_guide \,((20, 7, 2004), (20, 30), ABC)) =_{Program}$$
$$(\text{``The Bill''}, \text{````''}, 50, [Drama], M,$$
$$\text{``}Sun \ Hill \ continues \ to \ work \ at \ breaking \ the \ people \ smuggling \ operation\text{''})).$$

This fact states that the program on 20 July 2004 at 8.30pm on channel ABC has title "The Bill", no subtitle, a duration of 50 minutes, genre drama, a classification for mature audiences, and synopsis "Sun Hill continues to work at breaking the people smuggling operation".

There are a number of simple subsidiary functions that are defined as follows. Note that the definition of the function add is in the base of common beliefs and hence uses the modality $\Box$.

$$proj_{Occurrence} : State \rightarrow Occurrence$$

$$\Box_t \,\forall_{Occurrence} o. \,\forall_{Status} s.$$
$$((proj_{Occurrence} \,(o, s)) =_{Occurrence} o).$$

$$period : Occurrence \rightarrow Date \times Time \times Time$$

$$\Box_t \,\forall_{Date} d. \,\forall_{Time} t. \,\forall_{Channel} c.$$
$$((period \,(d, t, c)) =_{Date \times Time \times Time}$$
$$(d, t, (add \,(t, (proj_{Duration} \,(tv\_guide \,(d, t, c)))))))).$$

$$add : Time \times Duration \rightarrow Time$$

$$\Box \,\forall_{Hour} h. \,\forall_{Minute} m. \,\forall_{Duration} d.$$
$$((add \,((h, m), d)) =_{Time}$$
$$((60 \times h + m + d) \ div \ 60, (60 \times h + m + d) \ mod \ 60)).$$

$$proj_{Duration} : Program \rightarrow Duration$$
$$\Box_t \, \forall_{Title} t. \, \forall_{Subtitle} t'. \, \forall_{Duration} d. \, \forall_{(List \; Genre)} g. \, \forall_{Classification} c. \, \forall_{Synopsis} s.$$
$$((proj_{Duration} \; (t, t', d, g, c, s)) =_{Duration} d).$$

The belief base of the TV recommender includes the function *user_tv_time_acceptable* which has a definition that is obtained by belief acquisition and would look something like the following.

$$user\_tv\_time\_acceptable : Date \times Time \times Time \rightarrow \Omega$$
$$\Box_t \, \forall_{Date} d. \, \forall_{Time} t. \, \forall_{Time} t'.$$
$$((user\_tv\_time\_acceptable \; (d, t, t')) =_{\Omega}$$
$$if \; (weekday \; d) \wedge ((proj_{Hour} \; t) \geq 20) \wedge ((proj_{Hour} \; t') \leq 23) \; then \; \top$$
$$else \; if \; \neg(weekday \; d) \wedge ((proj_{Hour} \; t) \geq 12) \wedge ((proj_{Hour} \; t') \leq 25) \; then \; \top$$
$$else \; \bot).$$

This definition states that the user is willing to watch programs that start between 8pm and 11pm on weekdays and between midday and 1am on weekends. This information is obtained rather directly from the user by the belief acquisition algorithm.

The belief base of the TV recommender also includes the function *user_likes_tv_program* which has a definition that is obtained by belief acquisition (in this case, machine learning since the function learned generalises) and would look something like the following.

$$user\_likes\_tv\_program : Program \rightarrow \Omega$$
$$\Box_t \, \forall_{Program} x.$$
$$((user\_likes\_tv\_program \; x) =_{\Omega}$$
$$if \; (proj_{Title} \circ (=_{Title} \; \text{``NFL Football''}) \; x) \; then \; \top$$
$$else \; if \; (proj_{(List \; Genre)} \circ (listExists_1 \; genre \circ (< 0)) \; x) \; then \; \bot$$
$$else \; if \; (proj_{Title} \circ StringToText \circ (listExists_1 \; (=_{String} \; \text{``sport''})) \; x) \; then \; \top$$
$$else \; if \; (proj_{(List \; Genre)} \circ (listExists_1 \; (=_{Genre} \; Current\_Affairs)) \; x) \; then \; \bot$$
$$else \; if \; (proj_{Title} \circ (=_{Title} \; \text{``2004 ICC Championship Trophy''}) \; x) \; then \; \top$$
$$else \; if \; (proj_{Synopsis} \circ StringToText \circ (listExists_1 \; (=_{String} \; \text{``american''})) \; x) \; then \; \top$$
$$\vdots$$
$$else \; \bot).$$

Much more detail on how this function is learned is contained in [2].

Another subsidiary function needed is *user_diary_free* that has signature

$$user\_diary\_free : Date \times Time \times Time \rightarrow \Omega.$$

The definition of this function is in the belief base of the diary agent. The TV agent has access to this definition because of an interaction axiom that makes the belief base of the diary agent accessible to the TV agent.

We can now define the evidence variables.

The first evidence variable is

$$ultp : State \rightarrow \Omega$$
$$\Box_t \forall_{State} x.$$
$$((ultp \; x) =_\Omega$$
$$(proj_{Occurrence} \circ tv\_guide \circ user\_likes\_tv\_program \; x)).$$

This evidence variable is intended to indicate whether or not the user likes a program (irrespective of whether or not the user would be able to watch it).

The second evidence variable is

$$udiary : State \rightarrow \Omega$$
$$\Box_t \forall_{State} x.$$
$$((udiary \; x) =_\Omega$$
$$(proj_{Occurrence} \circ period \circ user\_diary\_free \; x)).$$

This evidence variable is intended to indicate whether or not the user's diary is free during the time that the program is on.

The third evidence variable is

$$uaccept : State \rightarrow \Omega$$
$$\Box_t \forall_{State} x.$$
$$((uaccept \; x) =_\Omega$$
$$(proj_{Occurrence} \circ period \circ user\_tv\_time\_acceptable \; x)).$$

This evidence variable is intended to indicate whether or not the user is willing to watch television at the time the program is on.

The intuition is that the three evidence variables are features that together can reasonably be used by the TV agent to make recommendations for the user.

Next we turn to the result variables.

The first result variable is the function

$$user : State \rightarrow \Omega$$

which models the user as an oracle about occurrences of programs. Given a state, *user* returns $\top$ if the user intends to watch the program whose occurrence is in the first component of the state; otherwise, *user* returns $\bot$. Of course, the definition of *user* is not available to the agent. But it can observe values for this function by asking the user.

The second result variable is

$$recomm : State \rightarrow Status$$
$$\Box_t \, \forall_{Occurrence} o. \, \forall_{Status} s.$$
$$((recomm \, (o, s)) =_{Status} s).$$

The function *recomm* simply projects onto the status component of the state.

Figure 4 illustrates the evidence and result variables for the TV recommender. The corresponding influence diagram is given in Figure 5.
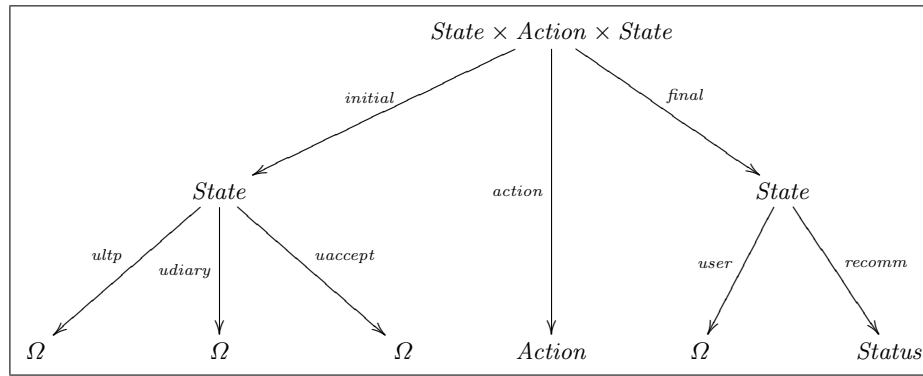
Fig. 4. Evidence and result variables for the TV recommender

Given in Figure 6 is a set of transition probability distributions and utilities that could reasonably arise in practice. (The utilities are given by the user.) The last column contains the expected utility of the action given the evidence for each row. In Figure 6 it is the case that the definition acquired by the agent for *uaccept* is actually not all that accurate and the user is willing to watch programs outside the periods given in the definition of this function. Secondly, it is assumed that the user is reasonably happy with states in which the agent recommends programs that they do not actually want to watch and therefore gives this state the utility value of 0.4. The policy corresponding to the situation is given in Figure 7 and its definition in the logic is as follows.

$$policy : State \rightarrow Action$$
$$\Box_t \, \forall_{State} s.$$
$$((policy \, s) =_{Action}$$
$$if \, (((ultp \, s) =_\Omega \top) \wedge ((udiary \, s) =_\Omega \top)) \, then \, RecommendYes$$
$$else \, RecommendNo).$$
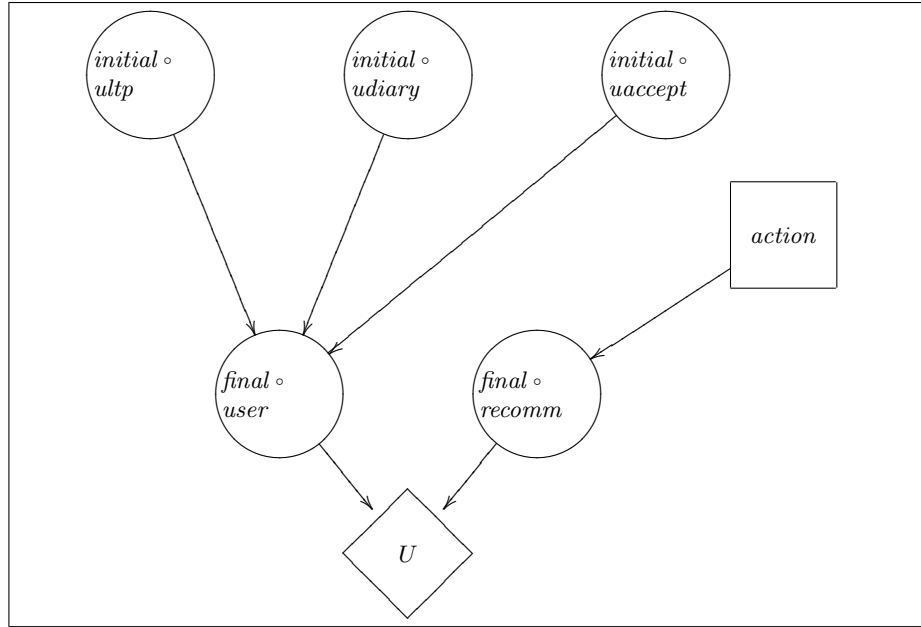
Fig. 5. Influence diagram for the TV recommender

$initial \circ (ultp, udiary, uaccept)$	$action$	$final \circ (user, recomm)$				
		$(\top, Yes)$	$(\top, No)$	$(\bot, Yes)$	$(\bot, No)$	
$(\top, \top, \top)$	$RecommendYes$	0.8	0.0	0.2	0.0	0.88
$(\top, \top, \top)$	$RecommendNo$	0.0	0.8	0.0	0.2	0.20
$(\top, \top, \bot)$	$RecommendYes$	0.4	0.0	0.6	0.0	0.64
$(\top, \top, \bot)$	$RecommendNo$	0.0	0.4	0.0	0.6	0.60
$(\top, \bot, \top)$	$RecommendYes$	0.0	0.0	1.0	0.0	0.40
$(\top, \bot, \top)$	$RecommendNo$	0.0	0.0	0.0	1.0	1.00
$(\top, \bot, \bot)$	$RecommendYes$	0.0	0.0	1.0	0.0	0.40
$(\top, \bot, \bot)$	$RecommendNo$	0.0	0.0	0.0	1.0	1.00
$(\bot, \top, \top)$	$RecommendYes$	0.1	0.0	0.9	0.0	0.46
$(\bot, \top, \top)$	$RecommendNo$	0.0	0.1	0.0	0.9	0.90
$(\bot, \top, \bot)$	$RecommendYes$	0.0	0.0	1.0	0.0	0.40
$(\bot, \top, \bot)$	$RecommendNo$	0.0	0.0	0.0	1.0	1.00
$(\bot, \bot, \top)$	$RecommendYes$	0.0	0.0	1.0	0.0	0.40
$(\bot, \bot, \top)$	$RecommendNo$	0.0	0.0	0.0	1.0	1.00
$(\bot, \bot, \bot)$	$RecommendYes$	0.0	0.0	1.0	0.0	0.40
$(\bot, \bot, \bot)$	$RecommendNo$	0.0	0.0	0.0	1.0	1.00
		1	0	0.4	1	

Fig. 6. Influence diagram conditioned on the evidence and action variables for the TV recommender

$initial \circ (ultp, udiary, uaccept)$	$action$
$(\top, \top, \top)$	$Recommend\,Yes$
$(\top, \top, \bot)$	$Recommend\,Yes$
$(\top, \bot, \top)$	$Recommend\,No$
$(\top, \bot, \bot)$	$Recommend\,No$
$(\bot, \top, \top)$	$Recommend\,No$
$(\bot, \top, \bot)$	$Recommend\,No$
$(\bot, \bot, \top)$	$Recommend\,No$
$(\bot, \bot, \bot)$	$Recommend\,No$

Fig. 7. Policy for the TV recommender corresponding to Figure 6

4   Conclusions

While the outlines of the architecture are now clear, much work remains to be done. For example, the belief acquisition algorithm is not yet fully implemented and investigated; the issue of how changing evidence and result variables affect the conditional probability tables in the Bayesian network needs further investigation; and further challenging applications for this approach to designing and building agents are needed.

Acknowledgements

References

1. C. Boutilier, T. Dean, and S. Hanks. Decision theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.
2. J.J. Cole, M.J. Gray, J.W. Lloyd, and K.S. Ng. Personalisation for user agents. In *Fourth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 05)*, 2005.
3. R.G. Cowell, A.P. Dawid, S.L. Lauritzen, and D.J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Statistics for Engineering and Information Sciences. Springer, 1999.
4. R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
5. J.W. Lloyd. *Logic for Learning*. Cognitive Technologies. Springer, 2003.
6. J.W. Lloyd. Modal higher-order logic for agents. Submitted for publication, 2004.
7. J.W. Lloyd. A method for belief acquisition. In preparation, 2005.
8. J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
9. R. Rivest. Learning decision lists. *Machine Learning*, 2(3):229–246, 1987.
10. S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, second edition, 2002.