Knowledge and Observations in the Situation Calculus

Ryan F. Kelly and Adrian R. Pearce
NICTA Victoria Research Laboratory
Department of Computer Science and Software Engineering
The University of Melbourne
Victoria, 3010, Australia
{rfk,adrian}@csse.unimelb.edu.au

ABSTRACT

We present a powerful new account of multi-agent knowledge in the situation calculus and an effective reasoning procedure for handling knowledge queries. Our approach generalizes existing work by reifying the observations made by each agent as the world evolves, allowing for agents that are partially or completely unaware of some of the actions that have occurred. This also enables agents to reason effectively about knowledge using only their internal history of observations, rather than requiring a full history of the world. The result is a more robust and flexible account of knowledge suitable for use in partially-observable multi-agent domains.

Categories and Subject Descriptors

I.2.4 [**Artificial Intelligence**]: Knowledge Representation Formalisms and Methods

General Terms

Theory

Keywords

Situation Calculus, Knowledge, Action, Observability

1. INTRODUCTION

The situation calculus [3], extended with knowledge [4] and concurrent actions [5], provides a rich formalism for modeling complex domains such as multi-agent systems. However, existing techniques for effective reasoning require an omniscient viewpoint, with queries posed relative to the current situation. This makes it difficult for situated agents to reason about their environment unless they have a complete history of the world, which is unrealistic in many domains.

We overcome this limitation by explicitly reifying the observations made by each agent as the world evolves. Knowledge is axiomatised such that each agent considers possible any situation compatible with what it has observed. We then extend the traditional definition of regression over knowledge queries to operate using the agent's history of observations rather than a full situation term. Our approach thus allows a situated agent to reason effectively about its own knowledge of the world even when it is only partially aware of the actions that have occurred.

2. BACKGROUND

Our work utilizes the situation calculus [3] enriched with concurrent actions [5] and multiple agents. A *basic action theory* is represented by $\mathcal{D}$. We utilize the notion of *action description predicates* as in [1], which include the familiar action possibility predicate $Poss$. The ordering relation $s <_\alpha s'$ should be read as "s' is in the future of s, and all intermediate actions satisfy α". Situations satisfying $S_0 \leq_{Poss} s$ are termed "legal situations" and are identified using $Legal(s)$. The *uniform formulae* [3] represent properties of the state of the world. We write ϕ for an arbitrary uniform formula and $\phi[s]$ for a uniform formula with its situation term replaced by s.

The standard semantics of knowledge [4, 5] are based on a "possible worlds" formulation. A knowledge fluent $K(agt, s', s)$ indicates that "in situation s, the agent agt considers the alternate situation s' to be possible". The macro **Knows** then acts as a shorthand for the standard possible-worlds definition of knowledge:

$$\mathbf{Knows}(agt, \phi, s) \stackrel{\text{def}}{=} \forall s' \,.\, K(agt, s', s) \rightarrow \phi[s']$$

The successor state axiom for the knowledge fluent (with multiple agents and concurrent actions) is:

$$\begin{aligned} &K(agt, s'', do(c, s)) \equiv \\ &\quad \exists s' \,.\, s'' = do(c, s') \wedge K(agt, s', s) \wedge Poss(c, s') \\ &\quad \wedge \forall a \in c. \left[agent(a) = agt \rightarrow SR(a, s) = SR(a, s')\right] \end{aligned}$$

It is also necessary to permit alternate initial situations, which are identified by $Init(s)$. Every situation is *rooted at* some initial situation, identified using $Root(s)$.

While powerful, this formulation has an important limitation: each agent is assumed to be aware of all actions that have occurred. While this is clearly infeasible for many multi-agent domains, it permits the regression operator [3] to be used for effective reasoning about knowledge queries. For formulations that do not make this assumption (see e.g. [2]) the successor state axiom for K involves universal quantification over situation terms and standard effective reasoning techniques cannot be applied.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

978-81-904262-7-5 (RPS) ©2007 IFAAMAS

To construct a more general formalism while retaining regression as a reasoning tool, we utilise the recently developed notion of *persistence condition* [1]. The operator $\mathcal{P}_{\mathcal{D}}(\phi, \alpha)$ produces a uniform formula that will hold in s if and only if ϕ will remain true in all future situations brought about by actions satisfying α (we say ϕ *persists* under α):

$$\mathcal{D} \models \left(\forall s' . s \leq_\alpha s' \rightarrow \phi[s']\right) \equiv \mathcal{P}_{\mathcal{D}}(\phi, \alpha)[s]$$

3. NEW SEMANTICS OF KNOWLEDGE

To generalize the existing account of knowledge in a robust way, we introduce a distinction between *actions*, which cause changes to the state of the world, and *observations*, which cause an agent to become *aware* of some change in the state of the world.

Definition 1. An *observation* is a notification received by an agent that makes it aware of some change in the state of the world. When an agent receives such a notification, we say that it "observed" or "perceived" that observation.

For simplicity we assume that agents perceive observations in the same instant as the actions causing them. We make no commitment as to how these notifications are generated, preferring a clean delineation between the task of observing change and the dynamics of knowledge update.

Let us introduce an additional sort OBSERVATION to the situation calculus, for the moment without commitment towards what this sort will contain. We then introduce the function $Obs(agt, c, s) = o$, returning a set of observations, to mean "when the actions c are performed in situation s, agent agt will perceive the observations o".

The concept of an *observation history* follows naturally - it is the sequence of all observations made by an agent as the world has evolved. Using ϵ to represent the empty history, the function $ObsHist$ can be defined to give the observation history associated with a particular situation:

$$\begin{aligned}
&Init(s) \rightarrow ObsHist(agt, s) = \epsilon \\
&ObsHist(agt, do(c, s)) = h \equiv \exists o . Obs(agt, c, s) = o \\
&\qquad \wedge (o = \{\} \rightarrow h = ObsHist(agt, s)) \\
&\qquad \wedge (o \neq \{\} \rightarrow h = o \cdot ObsHist(agt, s))
\end{aligned} \quad (1)$$

There is a strong analogue between situations and observation histories. A situation represents a complete, global history of all the actions that have occurred in the world, while an observation history is an agent's local history of all the observations it has made. The situation is an omniscient view of the world, the observation history a local view.

It is a basic tenet of epistemic reasoning that an agent's knowledge must depend solely on its local history: the knowledge that it started out with combined with the observations it has made since then. Clearly, s' should be related to s if their root situations are K-related, they result in the same sequence of observations, and s' is legal:

$$\begin{aligned}
\mathcal{D} \models K(agt, s', s) &\equiv K(Root(s'), Root(s)) \wedge \\
&Legal(s') \wedge ObsHist(agt, s') = ObsHist(agt, s)
\end{aligned} \quad (2)$$

Unfortunately this formulation cannot be used directly in a basic action theory, as these require that fluent change be specified using successor state axioms. We must formulate an appropriate axiom for K which enforces equation (2).

We first introduce $PbU(agt, c, s)$ as a notational convenience for "possible but unobservable", indicating that the actions c are possible in s, but no observations would be perceived by the agent agt if they are performed:

$$PbU(agt, c, s) \equiv Poss(c, s) \wedge Obs(agt, c, s) = \{\} \quad (3)$$

By stating that $s \leq_{PbU(agt)} s'$ we assert that an agent would perceive no observations were the world to change from situation s to s'. If it considers s possible then it must also consider s' possible. Following this intuition, the successor state axiom below captures the desired dynamics of K:

$$\begin{aligned}
&K(agt, s'', do(c, s)) \equiv \\
&\quad \left[Obs(agt, c, s) = \{\} \rightarrow K(agt, s'', s) \right] \\
&\wedge \ \left[Obs(agt, c, s) \neq \{\} \rightarrow \right. \\
&\qquad \exists c', s' . Obs(agt, c', s') = Obs(agt, c, s) \\
&\wedge Poss(c', s') \wedge K(agt, s', s) \wedge do(c', s') \leq_{PbU(agt)} s'' \left. \right]
\end{aligned} \quad (4)$$

If c was totally unobservable, the agent's state of knowledge does not change. Otherwise, it considers possible any legal successor to a possible alternate situation s' that can be brought about by actions c' that result in identical observations. It also considers possible any future of such a situation in which is would perceive no more observations.

Since situations satisfying $S_0 \leq_{PbU(agt)} s$ holds must be K-related to S_0 for equation (2) to be satisfied, we cannot use K to encode the agent's initial knowledge. We introduce K_0 for this purpose as follows:

$$\begin{aligned}
&Init(s) \rightarrow \\
&\quad K(agt, s'', s) \equiv \exists s' . K_0(agt, s', s) \wedge s' \leq_{PbU(agt)} s''
\end{aligned} \quad (5)$$

This axiomatisation is enough to ensure that knowledge meets to requirements set out in equation (2).

THEOREM 1. *For any basic action theory $\mathcal{D}$ axiomatising knowledge according to equations (4,5):*

$$\begin{aligned}
\mathcal{D} \models K(agt, s', s) &\equiv K(Root(s'), Root(s)) \wedge \\
&Legal(s') \wedge ObsHist(agt, s') = ObsHist(agt, s)
\end{aligned}$$

PROOF. Straightforward, using equations (1,3,4,5). □

Using this new formulation, an agent's knowledge is completely decoupled from the global notion of actions, instead depending only on the local information that it has observed. By allowing different kinds of term in the OBSERVATION sort, and axiomatising the *Obs* function in different ways, a wide variety of domains can be modelled.

Let us begin by considering the standard account of knowledge from [4]. Its basic assumption that "all agents are aware of all actions" can be captured by allowing the OBSERVATION sort to contain ACTION terms and including the following as an axiom:

$$a \in Obs(agt, c, s) \equiv a \in c \quad (6)$$

What about sensing information? We can extend the OBSERVATIONS sort to contain terms of the form ($Action = Result$) and axiomatize like so:

$$\begin{aligned}
&(a = r) \in Obs(agt, c, s) \equiv \\
&\qquad a \in c \wedge SR(a, s) = r \wedge agent(a) = agt
\end{aligned} \quad (7)$$

Using equations (6,7) as axioms, our new account of knowledge behaves identically to the standard account. To generalize this for partial observability of actions we introduce a new action description predicate specifying when actions will be observed by agents: $CanObs(agt, a, s)$ indicates that agent agt would observe action a being performed in situation s. We can then formulate $Obs()$ according to:

$$a \in Obs(agt, c, s) \equiv a \in c \wedge CanObs(agt, a, s)$$

There is an additional assumption in the standard handling of sensing actions: only the agent performing a sensing action is aware of its result. We can lift this assumption by introducing an analogous predicate $CanSense(agt, a, s)$ to indicate when sensing information is available to an agent. We then include bare action terms in an agent's observations when it observes the action but not its result, and *(Action=Result)* terms when it also senses the result:

$$a \in Obs(agt, c, s) \equiv a \in c \\ \wedge CanObs(agt, a, s) \wedge \neg CanSense(agt, a, s)$$

$$(a = r) \in Obs(agt, c, s) \equiv a \in c \wedge SR(a, s) = r \\ \wedge CanObs(agt, a, s) \wedge CanSense(agt, a, s)$$

This allows one to explicitly axiomatize the conditions under which an agent will be aware of the occurrence of an action. Even more complex domains can be modelled by further refining the Obs function, without modifying the dynamics of knowledge update.

4. REASONING

Standard regression techniques cannot be applied to our formalism, since equation (4) uses $\leq_{PbU}$ and so universally quantifies over situations. We have developed an effective reasoning procedure by using the persistence condition meta-operator [1] to augment the regression techniques in [4].

Assuming that the knowledge fluent K appears only in the context of a **Knows** macro, we propose the following to replace the regression clause in [4] for **Knows**:

$$\mathcal{R}_{\mathcal{D}}(\mathbf{Knows}(agt, \phi, do(c, s))) = \\ \exists o\,.\,Obs(agt, c, s) = o \wedge [o = \{\} \rightarrow \mathbf{Knows}(agt, \phi, s)] \\ \wedge\ [o \neq \{\} \rightarrow \mathbf{Knows}(agt, \forall c'.\,Obs(agt, c', s) = o \\ \wedge Poss(c', s) \rightarrow \mathcal{R}_{\mathcal{D}}(\mathcal{P}_{\mathcal{D}}(\phi, PbU(agt))[do(c', s)]), s)] \quad (8)$$

As required, this reduces a knowledge query at $do(c, s)$ to a knowledge query at s. It is also intuitively appealing: to know that ϕ holds, the agent must know that in all situations that agree with its observations, ϕ cannot become false without it perceiving an observation.

We must also specify the regression of **Knows** in the initial situation, as equation (5) also uses $\leq_{PbU}$. This clause results in standard first-order modal reasoning over the K_0 relation, similarly to [4]:

$$\mathcal{R}_{\mathcal{D}}(\mathbf{Knows}(agt, \phi, S_0)) = \\ \forall s\, K_0(agt, s, S_0) \rightarrow \mathcal{P}_{\mathcal{D}}(\phi, PbU(agt))[s] \quad (9)$$

The proof that our modified regression operator in equations (8,9) preserves equivalence proceeds by expanding the definition for **Knows** using our new successor state axiom for K, collecting sub-formulae that match the form of the **Knows** macro, and using regression and persistence to render the resulting knowledge expression uniform in s. Space restrictions prohibit a detailed exposition.

We can thus handle knowledge queries using regression, the standard technique for effective reasoning in the situation calculus. However, it would be unreasonable for a situated agent to ask "do I know ϕ in the current situation?" using the situation calculus query $\mathcal{D} \models \mathbf{Knows}(agt, \phi, s)$, as it cannot be expected to have the full current situation s. However, it will have its current observation history h. The regression rules in equations (8,9) can be modified to operate using an observation history rather than a situation term, with the result being significantly simpler due to the absence of empty observation sets:

$$\mathcal{R}_{\mathcal{D}}(\mathbf{Knows}(agt, \phi, o \cdot h)) = \\ \mathbf{Knows}(agt, \forall c\,.\,Obs(agt, c, s) = o \\ \wedge Poss(c, s) \rightarrow \mathcal{R}_{\mathcal{D}}(\mathcal{P}_{\mathcal{D}}(\phi, PbU(agt))[do(c, s)]), h) \quad (10)$$

$$\mathcal{R}_{\mathcal{D}}(\mathbf{Knows}(agt, \phi, \epsilon)) = \\ \forall s\,.\,K_0(agt, s, S_0) \rightarrow \mathcal{P}_{\mathcal{D}}(\phi, PbU(agt))[s] \quad (11)$$

It is a straightforward consequence of Theorem 1 that $\mathcal{R}_{\mathcal{D}}(\mathbf{Knows}(agt, \phi, s))$ and $\mathcal{R}_{\mathcal{D}}(\mathbf{Knows}(agt, \phi, ObsHist(s)))$ are equivalent. Agents can thus reason about their own knowledge using only their local information.

5. CONCLUSIONS

In this paper we have significantly increased the scope of the situation calculus for modeling knowledge in complex domains, where there may be multiple agents and partial observability of actions. By explicitly reifying the observations made by each agent as the world evolves, we have generalized the dynamics of knowledge update. Despite requiring universal quantification over future situations, we have shown that the regression operator can be adapted to allow effective reasoning within our new formalism. It can also be used to reason from the internal perspective of a single agent, allowing agents to reason about their own world based solely on their local information. With our new semantics of knowledge, the situation calculus is well positioned for representing, reasoning about, and implementing more complex, realistic multi-agent systems.

6. REFERENCES

[1] R. F. Kelly and A. R. Pearce. Property persistence in the situation calculus. In *IJCAI-07: Proc. of the 20th Int. Joint Conf. on A.I.*, 2007.

[2] Y. Lespérance, H. J. Levesque, and R. Reiter. A situation calculus approach to modeling and programming agents. In R. A. and M. Wooldridge, editors, *Foundations and Theories of Rational Agency*, pages 275–299. Kluwer, 1999.

[3] F. Pirri and R. Reiter. Some contributions to the metatheory of the situation calculus. *Journal of the ACM*, 46(3):325–361, 1999.

[4] R. Scherl and H. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144:1–39, 2003.

[5] R. B. Scherl. Reasoning about the interaction of knowledge, time and concurrent actions in the situation calculus. In *Proc. of the 18th International Conference on Artificial Intelligence*, pages 1091–1098, 2003.