Artificial Intelligence 221 (2015) 1–35

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Asynchronous knowledge with hidden actions in the situation calculus

Ryan F. Kelly, Adrian R. Pearce *

Department of Computing and Information Systems, The University of Melbourne, Victoria, 3010, Australia

A R T I C L E I N F O

Article history:
Received 22 April 2013
Received in revised form 7 December 2014
Accepted 20 December 2014
Available online 29 December 2014

Keywords:
Situation calculus
Knowledge
Epistemic reasoning

A B S T R A C T

We present a powerful new account of multi-agent knowledge in the situation calculus and an automated reasoning procedure for knowledge queries. Existing accounts of epistemic reasoning in the situation calculus require that whenever an action occurs, all agents *know* that an action has occurred. This demands a level of synchronicity that is unreasonable in many multi-agent domains. In asynchronous domains, each agent's knowledge must instead account for arbitrarily-long sequences of *hidden actions*. By using a persistence condition meta-operator to augment traditional regression techniques, we show how agents can reason about their own 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 in the situation calculus suitable for asynchronous, partially-observable multi-agent domains.

1. Introduction

In their landmark paper *Knowledge, Action, and the Frame Problem*, Scherl and Levesque [45] incorporate knowledge-producing actions into the situation calculus [37], inheriting Reiter's solution to the frame problem [38] and so enabling use of the regression operator to reason about the changing knowledge of an agent. Extensions to multiple agents [48] and concurrent actions [44] have produced an expressive logic of knowledge, action and change in which regression provides an automated reasoning procedure

While powerful, this formalism has a restriction that can make it unsuitable for modelling complex multi-agent domains. It requires that whenever an action occurs, all agents *know* that an action has occurred, demanding a level of synchronicity that is unreasonable in many multi-agent domains. If this restriction is lifted then each agent's knowledge must account for arbitrarily-long sequences of hidden actions [21], and proofs about knowledge must use a second-order induction axiom for quantifying over all future situations. This precludes the use of regression for automated reasoning. It also makes it difficult for agents to reason about their own knowledge, as they may not have enough information to formulate an appropriate query.

We overcome this limitation by combining two elements – an explicit representation of an agent's local perspective and a persistence condition meta-operator – to formulate an account of knowledge in the situation calculus that can faithfully represent the hidden actions inherent in asynchronous domains while maintaining regression as a key tool for automated reasoning.

* Corresponding author. Tel.: +61 3 8344 1399; fax: +61 3 9348 1184.
E-mail addresses: ryan@rfk.id.au (R.F. Kelly), adrianrp@unimelb.edu.au (A.R. Pearce).

http://dx.doi.org/10.1016/j.artint.2014.12.005

To decouple knowledge from the specific actions that have occurred, we explicitly reify the local *observations* made by each agent, so that every situation corresponds to an agent-local *view*. An agent is said to know proposition ϕ if and only if ϕ is true in all situations matching its current view. Our work thus has strong parallels with the classic view-based account of knowledge of Halpern and Moses [15], but grounded in the situation calculus and with an emphasis on regression for automated reasoning. This work appeared in preliminary form in [17]; the current paper presents an expanded treatment including proofs, properties and application domains.

The main challenge we overcome is developing regression rules that can handle arbitrarily long sequences of hidden actions. To ensure its knowledge is valid, an agent must reason about all future situations that are compatible with its observations. Such universal quantification over situation terms requires a second-order induction axiom, which the standard regression operator cannot handle. In previous work we have developed the *persistence condition* operator to handle this induction as a meta-level fixpoint calculation [19]. Using this operator to augment the regression techniques developed by Scherl and Levesque [45], we maintain their elegant solution to the frame problem while handling arbitrarily long sequences of hidden actions.

The formulation is shown to respect basic intuitions about how knowledge should behave, and to preserve important properties of the agent's epistemic state through the occurrence of actions. Moreover, it is *elaboration tolerant*, automatically preserving these properties in the face of more complex information-producing actions, such as the guarded sensing actions of [32], that can be tricky to axiomatize correctly in existing formalisms.

Decoupling knowledge from action in this manner makes it easy to express varying degrees of observability, from actions that are public through to ones that are completely hidden. To illustrate, we present a running example based on the well-known "Hunt the Wumpus" domain [41] which has been previously used to demonstrate the interleaved sensing and action typical of realistic domains in the situation calculus [42]. Our variant is complicated by the presence of multiple agents and partial observability of actions:

> Ann and Bob are hunting a Wumpus in a dungeon with many interconnecting rooms. They can fully observe each other's actions if they are in the same room, can hear each other's actions from adjacent rooms, and have no other means of synchronisation.
> Like any Wumpus, this one does not move, causes a stench in all adjacent rooms, and if shot will emit a piercing scream that can be heard anywhere in the dungeon.
> Can Ann and Bob coordinate their knowledge and actions in order to find and shoot the Wumpus?

The possibility of hidden actions makes this domain difficult to represent, let alone reason about, in standard theories of knowledge in the situation calculus. Our approach offers a straightforward formulation and an automated reasoning procedure.

Further demonstrating the utility of our approach, we show how the new regression rules can be applied using an agent's individual view, rather than requiring a full situation term. Agents can thus use our techniques to reason about their own knowledge using only their local information, making the formalism suitable both for reasoning *about*, and for reasoning *in*, rich multi-agent domains.

The end result is a significantly more general and robust theory of knowledge in the situation calculus that still permits an automated reasoning procedure. There is a large body of work that could benefit from our formalism, including: specification and verification of multi-agent systems [52]; theories of coordination [13] and ability [22]; reasoning about the epistemic feasibility of plans [23]; analysing multi-player games [1]; and our own work on the computation of complex epistemic modalities [18] and the cooperative execution of Golog programs [16].

The paper now proceeds with a review of the standard account of multi-agent knowledge in the situation calculus, before treating the individual knowledge of each agent in the face of hidden actions. We develop the axioms for our new observation-based account of knowledge in Section 3 and develop a regression rule for our formalism using the persistence condition operator in Section 4. Potential applications are shown in Section 5, where we show how our approach to axiomatising observations overcomes several difficulties encountered in previous formulations. An illustrative example of its use for reasoning about a partially-observable domain is provided in Section 6.

2. Background

Our work utilises the situation calculus [28,37] with multiple agents [48] and concurrent actions [40], and we begin from the standard account of knowledge due to Scherl and Levesque [45]. Several conservative extensions to the standard situation calculus meta-theory are also employed: the *Poss* predicate is subsumed by a general class of *action description predicates*; the unique names axioms for actions are subsumed by a general *background theory*; reasoning is performed using the *single-step regression operator* along with a new reasoning tool called the *persistence condition* operator [19].

There are of course a wide range of related formalisms for reasoning about knowledge, action and change, which we do not directly consider in this paper. We find the notation and meta-theory of the situation calculus particularly suitable for expressing our main ideas. Moreover, the strong underlying similarities between the major action formalisms should allow these ideas to transcend the specifics of the situation calculus [46,53,54].