

An On-line Decision-Theoretic Golog Interpreter.

Mikhail Soutchanski

Dept. of Computer Science
University of Toronto
Toronto, ON M5S 3H5
mes@cs.toronto.edu

Abstract

We consider an on-line decision-theoretic interpreter and incremental execution of Golog programs. This new interpreter is intended to overcome some limitations of the off-line interpreter proposed in [Boutilier *et al.*, 2000]. We introduce two new search control operators that can be mentioned in Golog programs: the on-line interpreter takes advantage of one of them to save computational efforts. In addition to sensing actions designed to identify outcomes of stochastic actions, we consider a new representation for sensing actions that may return both binary and real valued data at the run time. Programmers may use sensing actions explicitly in Golog programs whenever results of sensing are required to evaluate tests. The representation for sensing actions that we introduce allows the use of regression, a computationally efficient mechanism for evaluation of tests. We describe an implementation of the on-line incremental decision-theoretic Golog interpreter in Prolog. The implementation was successfully tested on the B21 robot manufactured by RWI.

1 Introduction

The aim of this paper is to provide a new on-line architecture of designing controllers for autonomous agents. Specifically, we explore controllers for mobile robots programmed in DT-Golog, an extension of the high-level programming language Golog. DTGolog aims to provide a seamless integration of a decision-theoretic planner based on Markov decision processes with an expressive set of program control structures available in Golog. The motivation for this research is provided in [Boutilier *et al.*, 2000], and we ask the reader to consult that paper for additional technical details and arguments why neither model-based planning, nor programming alone can manage the conceptual complexity of devising controllers for mobile robots. The DTGolog interpreter described in [Boutilier *et al.*, 2000] is an off-line interpreter that computes the optimal conditional policy π , the probability pr that π can be executed successfully and the expected utility u of the policy. The semantics of a DTGolog program p is defined by the predicate $BestDo(p, s, h, \pi, u, pr)$, where s is a starting situation and h is a given finite horizon. The policy π returned by the off-line interpreter is a Golog program consisting of the sequential composition of agent actions, $senseEffect(a)$ sensing

actions (which serve to identify a real outcome of a stochastic action a) and conditionals (**if** ϕ **then** π_1 **else** π_2), where ϕ is a situation calculus formula that provides a test condition to decide which branch of a policy the agent should take given the result of sensing. The interpreter looks ahead up to the last choice point in each branch of the program (say, between alternative primitive actions), makes the choice and then proceeds backwards recursively, deciding at each choice point what branch is optimal. It is assumed that once the off-line DTGolog interpreter has computed an optimal policy, the policy is given to the robotics software to control a real mobile robot.

However this off-line architecture has a number of limitations. Imagine that we are interested in executing a program $(p_1; p_2)$ where both sub-programs p_1 and p_2 are very large nondeterministic DTGolog programs designed to solve ‘independent’ decision-theoretic problems: p_2 is supposed to start in a certain set of states, but from the perspective of p_2 it is not important what policy will be used to reach one of those states. Intuitively, we are interested in computing first an optimal policy π_1 (that corresponds to p_1), executing π_1 in the real world and then computing and executing an optimal policy π_2 . But the off-line interpreter can return only the optimal policy π that corresponds to the whole program $(p_1; p_2)$, spending on this computation more time than necessary: to compute and execute π_2 it is not relevant what decisions have to be made during the execution of π_1 . The second limitation becomes evident if in addition to $senseEffect$ sensing actions serving to identify outcomes of stochastic actions, we need to explicitly include sensing actions in the program (and those sensing actions cannot be characterized as stochastic actions with a fixed finite set of outcomes). Imagine that we are given a program

$p_1; (\pi v, t). [(now(t))?; sense(Q, v, t); \mathbf{if} \phi \mathbf{then} p_2 \mathbf{else} p_3]$, where $sense(Q, v, t)$ is a sensing action that returns at time t a measurement v of a quantity Q (e.g., the current coordinates, the battery voltage) and the condition ϕ depends on the real data v that will be returned by sensors (in the program above p_1, p_2, p_3 are sub-programs, the choice operator π binds variables v and t and the test $now(t)?$ grounds the current time). Intuitively, we want to compute an optimal policy π_1 (corresponding to p_1) off-line, execute π_1 in the real world, sense v and then compute and execute an optimal policy π_2 that corresponds to the conditional Golog program in brackets. But the off-line interpreter is not able to compute an optimal policy if a given program includes explicit sensing actions and no information is available about the possible results of sensing.

We propose to compute and execute optimal policies on-line using a new incremental decision theoretic interpreter. It

works in a step-by-step mode. Given a Golog program p and a starting situation s , it computes an optimal policy π and the program p' that remains when a first action a in π will be executed. At the next stage, the action a is executed in the real world. The interpreter gets sensory information to identify which outcome of a has actually occurred if a is a stochastic action: this may require doing a sequence of sensing actions on-line. The action a (and possibly sensing actions performed after that) results in a new situation. Then the cycle of computing an optimal policy and remaining program, executing the first action and getting sensory information (if necessary) repeats until the program terminates or execution fails. In the context of the incremental on-line execution, we define a new programming operator $optimize(p)$ that limits the scope of search performed by the interpreter. If p is the whole program, then no computational efforts are saved when the interpreter computes an optimal policy from p , but if the programmer writes $optimize(p_1); p_2$, then the on-line incremental interpreter will compute and execute step-by-step the Golog program p_1 without looking ahead to decisions that need to be made in p_2 . If the programmer knows that the sensing action $sense(Q, v, t)$ is necessary to evaluate the condition ϕ , then using the program

```
optimize(p1); ( $\pi t, v$ ).[(now(t))?.optimize(sense(Q, v, t));
if  $\phi$  then p2 else p3]
```

the required information about Q will be obtained before the incremental interpreter will proceed to the execution of the conditional.¹

Thus, the incremental interpretation of the decision-theoretic Golog programs needs an account of sensing (formulated in the situation calculus) that will satisfy several criteria which naturally arise in the robotics context.

- Sensing actions have to accommodate both binary and real valued data returned from sensors.
- There has to exist a computationally efficient mechanism of using both sensing information and the information about the initial situation given the specifications of how a dynamical system evolves from one situation to another. Regression [Reiter, 2000] proved to be an efficient approach to the solution of the forward projection task (the evaluation of situation calculus formulas with a ground situation term), and for this reason it is desirable to integrate regression with reasoning about sensory information.
- Certain properties in the world vary unpredictably in time and cannot be modeled. Future sensing actions have to allow updating of the currently available information about unmodeled properties, but updates should not lead to inconsistencies.

¹The following simple scenario provides an illustration. Imagine an agent that has an airline ticket from Rome to Seattle and wants to compute an optimal policy of traveling from Rome to a conference venue in Seattle. This computational task consists of several independent decision theoretic tasks: compute an optimal policy of traveling from home in Rome to a local international airport and then execute it, sense the gate from where the flight will depart, fly to an international airport in Seattle and after arrival compute an optimal policy of traveling from the airport to the conference venue. A Golog programmer does not need any information about the number of gates in Rome and about assignment of flights to gates to write a program providing suitable constraints on the search for a globally optimal policy.

- Ideally, the specification of sensing actions in the situation calculus has to lead to a natural and sound implementation, e.g., in Prolog.

There are several accounts of sensing in the situation calculus that address different aspects of reasoning about sensory and physical actions [Bacchus *et al.*, 1995; De Giacomo and Levesque, 1999a; 1999b; Funge, 1998; Grosskreutz, 2000; Lakemeyer, 1999; Levesque, 1996; Pirri and Finzi, 1999; Scherl and Levesque, 1993].

Below we propose a new representation of sensing that simplifies reasoning about results of sensing, does not require consistency of sensory information with the domain theory, leads to a natural and sound implementation and has connections with representation of knowledge and sensing considered in [Reiter, 2000].

In Section 2 we recall the representation of the decision theoretic domain introduced in [Boutilier *et al.*, 2000]. In Section 3 we propose a representation for sensing actions and consider several examples. In section 4 we consider the on-line incremental decision-theoretic interpreter. Section 6 discusses connections with previously published papers.

2 The decision theoretic problem representation

The paper [Boutilier *et al.*, 2000] introduces the representation of problem domains that do not include sensing actions. The representation is based on the distinction between agent actions (which can be either deterministic or stochastic) and nature's actions which correspond to separate outcomes of a stochastic agent action. Nature's actions are considered deterministic. They cannot be executed by the agent itself, therefore they never occur in policies which the agent executes. When the agent does a stochastic action a in a situation s , nature chooses one of the outcomes n of that action and the situation $do(n, s)$ is considered as one of resulting situations. In accordance with this perspective, the evolution of a stochastic transition system is specified by precondition and successor state axioms which never mention stochastic agent actions, but mention only deterministic agent actions and nature's actions. In [Boutilier *et al.*, 2000], it is suggested to characterize the DTGolog problem domain by: 1) the predicate $agentAction(a)$ which holds if a is an agent action, 2) the predicate $stochastic(a, s, n)$ meaning that nature's action n is one of outcomes of the stochastic agent action a in the situation s , 3) the function symbol $prob(n, s)$ that denotes the probability of nature's action n in situation s , and 4) the predicate $senseCond(n, \phi)$ specifying the test condition ϕ serving to identify the outcome n of the last stochastic action, 5) the function symbol $reward(do(a, s))$ denotes rewards and costs as functions of the current situation $do(a, s)$, the action a or both.

As an example, imagine a robot moving between different locations: the process of going is initiated by a deterministic action $startGo(l_1, l_2, t)$ but is terminated by a stochastic action $endGo(l_1, l_2, t)$ that may have two different outcomes: successful arrival $endGoS(l_1, l_2, t)$ and unsuccessful stop in a place different from the destination $endGoF(l_1, l_3, t)$ (the robot gets stuck in the hall or cannot enter an office because its door is closed). We represent the process of moving between locations l_1 and l_2 by the relational fluent $going(l_1, l_2, s)$ and represent a (symbolic) location of the robot by the relational fluent $robotLoc(l, s)$ meaning that l (the office of an employee, the hall or the main office) is the place where the robot

is. The transitions in the stochastic dynamical system describing the robot's motion are characterized by the following precondition and successor state axioms.

$$Poss(startGo(l_1, l_2, t), s) \equiv \neg(\exists l, l') going(l, l', s) \wedge robotLoc(l_1, s),$$

$$Poss(endGoS(l_1, l_2, t), s) \equiv going(l_1, l_2, s),$$

$$Poss(endGoF(l_1, l_2, t), s) \equiv \exists l'. going(l_1, l', s) \wedge l' \neq l_2,$$

$$going(l, l', do(a, s)) \equiv (\exists t) a = startGo(l, l', t) \vee$$

$$going(l, l', s) \wedge \neg(\exists t) a = endGoS(l, l', t) \vee$$

$$going(l, l', s) \wedge \neg(\exists t, l'') a = endGoF(l, l'', t).$$

The real outcome n of a stochastic agent action a can be identified only from sensory information. This information has to be obtained by executing sensing actions. In the following section we propose a new representation for sensing actions providing a seamless integration with the representation considered in this section.

3 Sensing actions

In contrast to physical actions, sensing actions do not change any properties of the external world, but return values measured by sensors. Despite this difference, it is convenient to treat both physical and sensing actions equally in successor state axioms. This approach is justifiable if fluents represent what the robot 'knows' about the world (see Figure 1). More specifically, let the high-level control module of the robot be provided with the internal logical model of the world (the set of precondition and successor state axioms) and the axiomatization of the initial situation. The programmer expresses in this axiomatization his (incomplete) knowledge of the initial situation and captures certain important effects of the robot's actions, but some other effects and actions of other agents may remain unmodelled. When the robot does an action in the real world (i.e., the high-level control module sends a command to effectors and effectors execute this command), it does not know directly and immediately what effects in reality this action will produce: the high-level control module may only compute expected effects using the internal logical model. Similarly, if the robot is informed about actions executed by external agents, the high-level control module may compute expected effects of those exogenous actions (if axioms account for them). Thus, from the point of view of the programmer who designed the axioms, the high-level control module maintains the set of beliefs that the robot has about the real world. This set of beliefs needs feedback from the real world because of possible contingencies, incompleteness of the initial information and unobserved or unmodelled actions executed by other agents. To gain the feedback, the high-level control module requests information from sensors and they return required data. We find it convenient to represent information gathered from sensors by an argument of the sensing action: a value of the argument will be instantiated at the run time when the sensing action will be executed. Then, all the high-level control module needs to know is the current situation (action log): expected effects of actions on fluents can be computed from the given sequence of physical and sensing actions.

In the sequel, we consider only deterministic sensing actions, but noisy sensing actions can be represented as well.

We suggest representing sensing actions by the functional symbol $sense(q, v, t)$ where q is what to sense, v is term representing a run time value returned by the sensors and t is time; the predicate $senseAction(a)$ is true whenever a is a

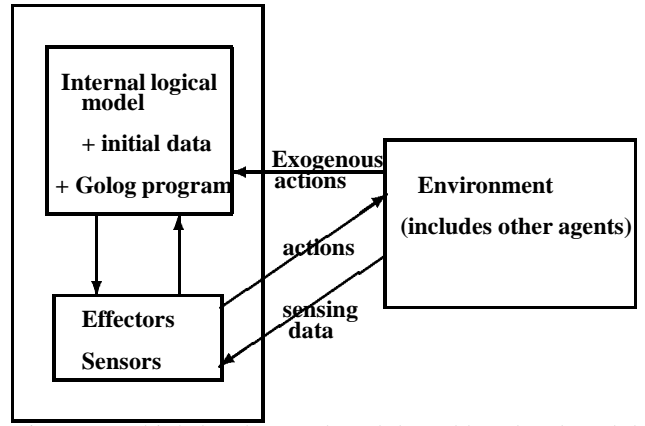


Figure 1: A high-level control module and low-level modules interacting with the environment.

sensing action. We proceed to consider several examples of successor state axioms that employ our representation.

1. Let $sense(Coord, l, t)$ be the sensing action that returns the pair $l = (x, y)$ of geometric coordinates of the current robot location on the two-dimensional grid and $gridLoc(x, y, s)$ be a relational fluent that is true if (x, y) are the coordinates of the robot's location in s . In this example we assume that all actions are deterministic (we do this only to simplify the exposition of this example). The process of moving (represented by the relational fluent $moving(x_1, y_1, x_2, y_2, s)$) from the grid location (x_1, y_1) to the point (x_2, y_2) is initiated by the deterministic instantaneous action $startMove(x_1, y_1, x_2, y_2, t)$ and is terminated by the deterministic action $endMove(x_1, y_1, x_2, y_2, t)$. The robot may also change its location if it is transported by an external agent from one place to another: the exogenous actions $take(x_1, y_1, t)$ and $put(x_2, y_2, t)$ account for this (and the fluent $transported(s)$ represents the process of moving the robot by an external agent). The following successor state and precondition axioms characterize aforementioned fluents and actions.

$$gridLoc(x, y, do(a, s)) \equiv$$

$$((\exists t, l) a = sense(Coord, l, t) \wedge xCrd(l) = x \wedge yCrd(l) = y) \vee$$

$$(\neg transported(s) \wedge (\exists t, x', y') a = endMove(x', y', x, y, t)) \vee$$

$$(transported(s) \wedge (\exists t) a = put(x, y, t)) \vee$$

$$gridLoc(x, y, s) \wedge \neg(\exists t, x', y') a = put(x', y', t) \wedge$$

$$\neg(\exists l, t) a = sense(Coord, l, t) \wedge$$

$$\neg(\exists x, y, x', y', t) a = endMove(x, y, x', y', t),$$

where $xCrd(l)$, $yCrd(l)$ denote, respectively, x and y components of the current robot location l .

$$moving(x_1, y_1, x_2, y_2, do(a, s)) \equiv$$

$$(\exists t) a = startMove(x_1, y_1, x_2, y_2, t) \vee$$

$$moving(x_1, y_1, x_2, y_2, s) \wedge$$

$$\neg(\exists t) a = endMove(x_1, y_1, x_2, y_2, t).$$

$$transported(do(a, s)) \equiv (\exists t, x, y) a = take(x, y, t) \vee$$

$$transported(s) \wedge \neg(\exists t, x', y') a = put(x', y', t) \vee$$

$$(\neg \exists x, y, x', y') moving(x, y, x', y') \wedge gridLoc(x, y, s) \wedge$$

$$(\exists t, l) a = sense(Coord, l, t) \wedge (xCrd(l) \neq x \vee yCrd(l) \neq y).$$

$$Poss(startMove(x_1, y_1, x_2, y_2, t), s) \equiv \neg transported(s) \wedge$$

$$\neg(\exists x, y, x', y') moving(x, y, x', y', s) \wedge gridLoc(x_1, y_1, s),$$

$$Poss(endMove(x_1, y_1, x_2, y_2, t), s) \equiv moving(x_1, y_1, x_2, y_2, s),$$

$$Poss(take(x, y, t), s) \equiv \neg transported(s) \wedge gridLoc(x, y, s) \wedge \neg(\exists x_1, y_1, x_2, y_2) moving(x_1, y_1, x_2, y_2, s),$$

$$Poss(put(x, y, t), s) \equiv transported(s).$$

Imagine that in the initial situation the robot stays at (0,0); later, at time 1, it starts moving from (0,0) to (2,3), but when the robot stops at time 11 and senses its coordinates at time 12, its sensors tell it that it is located at (4,4). The sensory information is inconsistent with the expected location of the robot, but this discrepancy can be attributed to unobserved actions of an external agent who transported the robot. Hence, the final situation is not

$$S_3 = do(sense(Coord, (4, 4), 12), do(endMove(0, 0, 2, 3, 11), do(startMove(0, 0, 2, 3, 1), S_0))),$$

as we originally expected, but the situation resulting from the execution of exogenous actions $take(2, 3, T_1)$ and $put(4, 4, T_2)$, in the situation when the robot ends moving, followed by the sensing action. The exogenous actions occurred at unknown times T_1, T_2 (we may say about the actual history only that $11 \leq T_1 < T_2 \leq 12$).²

2. The robot can determine its location using data from sonar and laser sensors. But if the last action was not sensing coordinates (x, y) , then the current location can be determined from the previous location and the actions that the robot has executed. The process of going from l to l' is initiated by the deterministic action $startGo(l, l', t)$ and is terminated by the stochastic action $endGo(l, l', t)$ (axiomatized in Section 2).

$$\begin{aligned} robotLoc(l, do(a, s)) &\equiv (\exists t, v, x, y) a = sense(Coord, v, t) \\ &\wedge xCrd(v) = x \wedge yCrd(v) = y \wedge \\ &((\exists p) l = office(p) \wedge inOffice(l, x, y) \vee \\ & l = Hall \wedge \neg(\exists p) inOffice(office(p), x, y)) \vee \\ &(\exists t, l_1) a = endGoS(l_1, l, t) \vee (\exists t, l_1) a = endGoF(l_1, l, t) \vee \\ &robotLoc(l, s) \wedge \neg(\exists t, l_2, l_3, v) (a = sense(Coord, v, t) \wedge \\ & a = endGoS(l_2, l_3, t) \wedge a = endGoF(l_2, l_3, t)). \end{aligned}$$

where the predicate $inOffice(l, x, y)$ is true if the pair (x, y) is inside of the office l , the functional symbols $bottomY(l)$, $topY(l)$, $rightX(l)$ and $leftX(l)$ represent coordinates of top left and bottom right corners of a rectangle that contains the office l inside.

$$\begin{aligned} inOffice(l, x, y) &\equiv (\exists y_b, y_t, x_l, x_r) \\ bottomY(l) &= y_b \wedge topY(l) = y_t \wedge \\ leftX(l) &= x_l \wedge rightX(l) = x_r \wedge \\ x_l \leq x \leq x_r \wedge y_b \leq y \leq y_t \end{aligned}$$

When the robot stops and senses coordinates, it determines its real location and the high-level control module can identify the outcome of the stochastic action $endGo(l, l', t)$: whether the robot stopped successfully or failed to arrive at the intended destination.

3. Let $give(item, person, t)$ be a stochastic action that has two different outcomes: $giveS(item, person, t)$ – the robot gives successfully an *item* to *person* at time t – and $giveF(item, person, t)$ – the action of giving an *item* to *person* is unsuccessful. Let $sense(Ackn, v, t)$ be the action of sensing whether delivery of the *item* to *person* was successful or not: if it was, then delivery is acknowledged and

²In the sequel, we do not consider how the discrepancy between results of a deterministic action and observations obtained from sensors can be explained by occurrences of unobserved exogenous actions. However, our example indicates that inconsistency between physical and sensory actions can be resolved by solving a corresponding diagnostic problem (e.g., see [McIlraith, 1998]).

$v=1$, if not – the result of sensing is $v=0$ (e.g., acknowledgment $v=1$ can be implemented if a recipient presses one of robot's buttons during a certain interval of time; if no button was pressed, then $v=0$). The following successor state axiom characterizes how the fluent $hasCoffee(person, s)$ changes from situation to situation: whenever the robot is in the office of *person* and it senses that one of its buttons was pressed, it is assumed that the *person* pressed a button to acknowledge that she has a coffee. From this sensory information, the high-level control module can identify whether the outcome of $give(item, person, t)$ was successful or not.

$$\begin{aligned} hasCoffee(pers, do(a, s)) &\equiv \\ &(\exists t) a = giveS(Coffee, pers, t) \vee hasCoffee(pers, s) \vee \\ &(\exists t, v) a = sense(Ackn, v, t) \wedge v = 1 \wedge \\ &robotLoc(office(pers), s). \end{aligned}$$

4. Assume that the robot has no devices to sense its location directly, but can measure its velocity (e.g., by counting revolutions of its motors); we model this by the sensing action $sense(Vel, v, t)$. For simplicity of presentation, the motion is uniform and the robot moves along the x axis; we also assume in this example that both $startMove$ and $endMove$ actions are deterministic (their precondition axioms are similar to axioms in the example 1). In this example, when the robot starts moving it moves with a certain velocity that remains constant but unknown, unless the robot measures its velocity.

Given the value returned by the velocity sensor, we are interested in computing the current location of the robot:³

$$\begin{aligned} moving(do(a, s)) &\equiv (\exists t) a = startMove(t) \vee \\ &moving(s) \wedge \neg(\exists t) a = endMove(t). \end{aligned}$$

$$\begin{aligned} velocity(v, do(a, s)) &\equiv (\exists t) a = sense(Vel, v, t) \vee \\ &(\exists t) a = endMove(t) \wedge v = 0 \vee \\ &(\neg \exists t) a = endMove(t) \wedge (\neg \exists t, v) a = sense(Vel, v, t) \wedge \\ &(\neg moving(s) \wedge v = 0 \vee moving(s) \wedge velocity(v, s)). \end{aligned}$$

$$\begin{aligned} locat(x, do(a, s)) &\equiv \\ &(\exists v, t_1, t_2, x') (a = sense(Vel, v, t_1) \wedge locat(x', s) \wedge \\ & start(s) = t_2 \wedge x = v \cdot (t_1 - t_2) + x') \vee \\ &(\exists v, t_1, t_2, x') (a = endMove(t_1) \wedge velocity(v, s) \wedge \\ & locat(x', s) \wedge start(s) = t_2 \wedge x = v \cdot (t_1 - t_2) + x') \vee \\ &(\neg \exists t) a = endMove(t) \wedge (\neg \exists t, v) a = sense(Vel, v, t) \wedge \\ &[\neg moving(s) \wedge locat(x, s) \vee moving(s) \wedge \\ & (\exists v, t_1, t_2, x') (velocity(v, s) \wedge locat(x', s) \wedge \\ & start(s) = t_2 \wedge time(a) = t_1 \wedge x = v \cdot (t_1 - t_2) + x')]. \end{aligned}$$

It is surprisingly straightforward to use regression in our setting to solve the forward projection task. Let \mathcal{D} be a background axiomatization of the domain (a set of successor state axioms, precondition axioms, unique name axioms and a set of first order sentences whose only situation term is S_0) and let $\phi(s)$ be a situation calculus formula that has the free variable s as the only situational argument. Suppose we are given a ground situational term S that may mention both physical and sensing actions. The forward projection task is to determine whether

$$\mathcal{D} \models \phi(S)$$

Regression is a computationally efficient way of solving the forward projection task [Reiter, 2000; Pirri and Reiter, 1999] when S does not mention any sensing actions. The representation introduced above allows us to use regression also in the

³We use the function symbol $start(s)$ to denote the starting time of the situation s and the function symbol $time(a)$ to denote the time when the action a occurs, see [Reiter, 1998] for details.

case when S mentions sensing actions explicitly. Thus, we can use regression to evaluate tests (ϕ)? in Golog programs with sensing actions: no modifications are required. In addition, our approach allows us to use an implementation in Prolog considered in [Reiter, 2000].

There is an interesting connection between our representation of ‘beliefs’ and sensing and the approach to representation of sensing and knowledge considered in Section 11.7 on knowledge-based programming [Reiter, 2000], where knowledge is defined using an epistemic fluent $K(s, \sigma)$. According to the main theorem of that section, knowledge of ϕ in situation σ (where σ is a ground situation term) is entailed by a basic action theory \mathcal{D} augmented by knowledge about the outcomes of all sense actions mentioned in σ iff regression of $\phi[\sigma]$ is entailed only by conjunction of unique names axioms and initial situation axioms augmented by regression of all sentences obtained as outcomes of sensing actions (the important condition when this theorem holds is that the theory of S_0 consists of conjunction of sentences declaring only what the agent knows about the world it inhabits, but there are no sentences declaring what is actually true of the world, or what the agent knows about what he knows). In other words, reasoning about knowledge can be reduced to provability of a sentence (whose only situation term is S_0) relative to axioms that either do not mention situation terms or mention only S_0 [Reiter, 2000]. Under the conditions when aforementioned theorem holds and when we consider sensing of fluents only, but not arbitrary sentences, one can prove that knowledge of ϕ in situation σ is entailed by conjunction of a basic action theory (where successor state axioms have occurrences of physical actions only) and axioms about knowledge of outcomes of sensing iff $\phi[\sigma]$ is entailed by our modification of a basic action theory, where successor state axioms mention both physical and sensory actions (see [Authors, forthcoming]). Intuitively, this is true because we replace conjunction of successor state axioms (that mention only physical actions) and axioms about outcomes of sensing by successor state axioms that mention both physical and sensory actions.

4 The incremental on-line DTGolog interpreter

The incremental DTGolog interpreter uses the predicate $IncrBestDo(p_1, s, p_2, h, \pi, u, pr)$ and the predicate $Final(p, s, \pi, u)$. The former predicate takes as input the Golog program p_1 , starting situation s , horizon h and returns the optimal conditional policy π , its expected utility u , the probability of success pr , and the program p_2 that remains after executing the first action from the policy π . The latter predicate tells when the execution of the policy completes: $Final(p, s, \pi, u)$ is true if either the program p is *Nil* (the null program) or *Stop* (a zero cost action that takes the agent to an absorbing state meaning that the execution failed), or if the policy π is *Nil* or *Stop*. In all these cases, u is simply the reward associated with the situation s . Note that in comparison to *BestDo*, *IncrBestDo* has an additional argument p_2 representing the program that remains to be executed.

$IncrBestDo(p_1, s, p_2, h, \pi, u, pr)$ is defined inductively on the structure of a Golog program p_1 . Below we consider its definition in the case when the program p_1 begins with a de-

terministic agent action.

$$\begin{aligned} IncrBestDo(a; p_1, s, p_2, h, \pi, u, pr) &\stackrel{def}{=} \\ &[\neg Poss(a, s) \wedge \\ & p_2 = Nil \wedge \pi = Stop \wedge pr = 0 \wedge u = reward(s) \vee \\ & Poss(a, s) \wedge p_2 = p_1 \wedge \exists(p', \pi', u', pr') \\ & IncrBestDo(p_2, do(a, s), p', h-1, \pi', u', pr') \wedge \\ & \pi = (a; \pi') \wedge u = reward(s) + u' \wedge pr = pr']. \end{aligned}$$

If a deterministic agent action a is possible in situation s , then we compute the optimal policy π' of the remaining program p_1 , its expected utility u' and the probability of successful termination pr' . Because a is a deterministic action, the probability pr' that the policy π' will complete successfully is the same as for the program itself; the expected utility u is a sum of a reward for being in s and the expected utility of continuation u' . If a is not possible, then the remaining program is *Nil*, and the policy is the *Stop* action, the expected utility is the reward in s and the probability of success is 0.

Other cases are defined similarly to *BestDo*, e.g., if the program begins with the finite nondeterministic choice ($\pi(x : \tau)p$); p' , where τ is the finite set $\{c_1, \dots, c_n\}$ and the choice binds all free occurrences of x in p to one of the elements:

$$\begin{aligned} IncrBestDo((\pi(x : \tau)p); p', s, pr, h, pol, u, pr) &\stackrel{def}{=} \\ IncrBestDo((p|_{c_1}^x \mid \dots \mid p|_{c_n}^x); p', s, pr, h, pol, u, pr) \end{aligned}$$

where $p|_c^x$ means substitution of c for all free occurrences of x in p . Thus, the optimal policy pol corresponds to the element c in τ that delivers the best execution. Note that the remaining program p_r is the same on the both sides of the definition.

Recall that policies are Golog programs as well. Moreover, if p is a Golog program that contains many nondeterministic choices, the optimal policy π computed from p is a conditional program that does not involve any nondeterministic choices. This observation suggests that programmers may wish to take advantage of the structure in a decision theoretic problem and use explicit search control operators that limit bounds where the search for an optimal policy has to concentrate. In addition to the standard set of Golog programming constructs, we introduce two new operators: $local(p)$ and $optimize(p)$. Intuitively, the program $local(p_1); p_2$ means the following. First, compute the optimal policy π_1 corresponding to the sub-program p_1 , then compute the optimal policy π corresponding to the program $\pi_1; p_2$. If both sub-programs p_1 and p_2 are highly nondeterministic, then using the operator $local(p_1)$ the programmer indicates where the computational efforts can be saved: there is no need in looking ahead further than p_1 to compute π_1 . Thus, in the case that a Golog program begins with $local(p)$

$$\begin{aligned} IncrBestDo(local(p_1); p_2, s, pr, h, \pi, u, pr) &\stackrel{def}{=} \\ (\exists p, \pi_1, u_1, pr_1) IncrBestDo(p_1; Nil, s, p, h, \pi_1, u_1, pr_1) \wedge \\ IncrBestDo(\pi_1; p_2, s, pr, h, \pi, u, pr). \end{aligned}$$

The construct $local(p_1)$ limits the search, but once the policy π_1 was computed and the first action in π_1 was executed, the remaining part of the program has no indication where the search may concentrate. For this reason, the programmer may find convenient to use another search control operator that once used persists in the remaining part of the program until the program inside the scopes of that operator terminates. This operator is called $optimize(p)$ and is specified by the following abbreviation.

$$\begin{aligned} \text{IncrBestDo}(\text{optimize}(p_1); p_2, s, p_r, h, \pi, u, pr) \stackrel{def}{=} \\ (\exists p') \text{IncrBestDo}(p_1; Nil, s, p', h, \pi, u, pr) \wedge \\ (p' \neq Nil \wedge p_r = (\text{optimize}(p'); p_2) \vee \\ p' = Nil \wedge p_r = p_2). \end{aligned}$$

According to this specification, an optimal policy π can be computed without looking ahead to the program p_2 ; hence, using $\text{optimize}(p_1)$ a programmer can express a domain specific procedural knowledge to save computational efforts. Note that when p' is Nil , i.e. there will be nothing in p_1 to execute after doing the only action in π , the remaining program p_r contains only p_2 .

4.1 Implementing the on-line interpreter

Given the definitions of *IncrBestDo* mentioned in the previous sub-section, we can consider now the on-line interpretation coupled with execution of Golog programs. The definitions of *IncrBestDo*($p_1, s, p_2, h, \pi, u, pr$) translate directly into Prolog clauses (we omit them here). The on-line interpreter calls the off-line *incrBestDo*($E, S, ER, H, Pol, U, U1, Probl$) interpreter to compute an optimal policy from the given program expression E , gets the first action of the optimal policy, commits to it, does it in the physical world, then repeats with the rest of the program. The following is such an interpreter implemented in Prolog:

```
online(E, S, H, Pol, U) :-
incrBestDo(E, S, ER, H, Poll, U1, Probl),
( final(ER, S, H, Poll, U1), Pol=Poll, U=U1 ;
  reward(R, S),
  Poll = (A : Rest),
  ( agentAction(A), not stochastic(A, S, L),
    doReally(A), /*execute A in reality*/
    !, /* commit to the result */
    online(ER, do(A, S), H, PolFut, UFut),
    Pol=(A : PolFut), U is R + UFut ;
    senseAction(A),
    doReally(A), /* do sensing */
    !, /*commit to results of sensing*/
    online(ER, do(A, S), H, PolFut, UFut),
    Pol=(A : PolFut), U is R + UFut ;
    agentAction(A), stochastic(A, S, L),
    doReally(A), /*execute A in reality*/
    !, /* commit to the result */
    senseEffect(A, S, SEff),
    diagnose(SEff, L, SN), /*What happened?*/
    online(ER, SN, H, PolFut, UFut),
    Pol=(A : PolF), U is R + UFut
  )
).
```

The on-line interpreter uses the Prolog cut (!) to prevent backtracking to the predicate *doReally*: we need this because once actions have been actually performed in the physical world, the robot cannot undo them.

In addition to predicates mentioned in section 2, the on-line interpreter uses the predicate *senseEffect*($A, S1, S2$), the predicate *diffSequence*(A, Seq) and the predicate *diagnose*($S2, OutcomesList, S3$). We describe below their meaning and show their implementation in Prolog.

Given the stochastic action A and the situation $S1$, the predicate *senseEffect*($A, S1, S2$) holds if $S2$ is the situation that results from doing a number of sensing actions necessary to differentiate between outcomes of the stochastic action A . The predicate *diffSequence*(A, Seq) holds if Seq is the sequence of sensing actions ($a_1; a_2; \dots; a_n$) specified by the programmer in the domain problem representation: this sequence is differentiating if after doing all actions in the sequence the action chosen by ‘nature’ as the outcome of stochastic action A can be uniquely identified.

```
senseEffect(A, S, SE) :- diffSequence(A, Seq),
  getSensorInput(S, Seq, SE).

getSensorInput(S, A, do(A, S)) :- senseAction(A),
  doReally(A). /*connect to sensors, get data
  for a free variable in A */

getSensorInput(S1, (A : Rest), SE) :-
  doReally(A), /* connect to sensors,
  get data */
  getSensorInput(do(A, S1), Rest, SE).
```

The predicate *diagnose*($S1, L, S2$) takes as its first argument the situation resulting from getting a sensory input: it contains ‘enough’ information to disambiguate between different possible outcomes of the last stochastic action A . The second argument is the list L of all outcomes that nature may choose if the agent executes the stochastic action A in $S1$ and the third argument is the situation that is the result of nature’s action which actually occurred. We can identify which action nature has chosen using the set of mutually exclusive test conditions *senseCond*(n_i, ϕ_i), where ϕ_i is a term representing a situation calculus formula: if ϕ_i holds in the current situation, then we know that nature has chosen the action n_i (n_i belongs to the list L).

```
diagnose(SE, [N], do(N, SE)) :-
  senseCond(N, C), holds(C, SE).

diagnose(SE, [N|Outcomes], SN) :- senseCond(N, C),
  ( holds(C, SE), SN=do(N, SE) ;
    not holds(C, SE),
    diagnose(SE, Outcomes, SN) ).
```

5 Examples

In this section we describe a simple, but realistic domain where natural constraints on robot behavior can be conveniently expressed in a Golog program. In domains of this type, we can use the on-line interpreter to find and execute incrementally an optimal policy.

Imagine a robot that has to deliver coffee from the main office where a coffee machine is located. A secretary can put a cup of coffee on the robot and we model this by saying that the robot executes the deterministic action *pickup*($item, t$). We model the process of moving by the deterministic action *startGo*(l_1, l_2, t) and stochastic action *endGo*(l_1, l_2, t). When the robot arrives at the door of an employee, it asks the employee to take delivered items; we model this by the stochastic action *give*($item, person, t$): if the employee is in her office and takes the item, then this action has a successful outcome *giveS*($item, person, t$); but if the employee is out, then *give*($item, person, t$) fails, reflected by outcome *giveF*($item, person, t$) (see example 4 in Section 3 for details). When the robot is in the main office, it can sense the battery voltage by doing the action *sense*($Battery, v, t$) and, if necessary, connect to the charger.

This domain is modeled using fluents such as *going*(l_1, l_2, s), *robotLoc*(l, s), *carrying*($item, s$), *hasCoffee*($person, s$), *voltage*(x, s), etc., with obvious intent. Several other predicates, function symbols and constants (e.g., names of people, items and locations) are also used.⁴ Some examples of preconditions and successor

⁴We use the predicate *wantsCoffee*(p, t_1, t_2) to assert that the person p wants coffee anytime in the interval $[t_1, t_2]$, the function symbol *travelTime*(l_1, l_2) to denote the travel time between locations l_1 and l_2 and the predicate *in*(x, y) to assert that x is an element of y .

state axioms have already been described in Section 2 and examples 3, 4 and 5 in Section 3. In addition to them, we need the successor state axiom that characterizes the relational fluent $status(p, st, s)$: a person p can be in or out of an office and the robot can determine the current status st when it gives an item:

$$\begin{aligned} status(person, st, do(a, s)) \equiv \\ (\exists t) a = giveS(item, person, t) \wedge st = In \vee \\ (\exists t) a = giveF(item, person, t) \wedge st = Out \vee \\ status(person, st, s) \wedge \\ \neg(\exists i, p, t) a = giveS(i, p, t) \wedge \neg(\exists i, p, t) a = giveF(i, p, t). \end{aligned}$$

Examples of sense conditions and axioms specifying probabilities are provided below: $prob(endGoS(l_1, l_2, t), s) = 0.99$, $senseCond(endGoS(l_1, l_2, t), robotLoc(l_2))$, $prob(giveS(Coffee, Craig, t), s) = 0.8$.

The robot receives a positive reward for delivering coffee in time. If a person wants a coffee from t_1 to t_2 , then the robot gets the highest reward $\frac{t_2 - t_1}{2}$ when it gives coffee at t_1 to that person, the reward declines linearly with time and at moments $\frac{t_1}{2}$ and t_2 the reward becomes zero; the reward is negative outside the interval $(\frac{t_1}{2}, t_2)$. Because all actions have a time argument, it is easy to compute the reward as the maximum of the linear function mentioned above given a sequence of actions that includes the action of giving coffee as the last action (inequalities between the time arguments of actions in the sequence provide a set of temporal constraints). The rewards for doing all other actions are zero.

In the initial situation, the robot is located in the main office and there are several requests for coffee. The MDP problem is to find an optimal policy that corresponds to delivering coffee in the optimal way.

The following compact Golog procedures express natural constraints of the robot's behavior (with abbreviations MO (main office), C (coffee)). The argument r of the procedure $deliver(r)$ denotes the finite range of offices and each office is the finite range of people who occupy it; for example, the argument of this procedure can be the finite set of offices $\{LP271, LP276, LP290b, LP269\}$, where the office $LP271$ is the finite set of names $\{Steve, Maurice, Yves\}$.

```

proc deliver( $r$ )
  while  $(\exists x) in(x, r) \wedge (\exists person) in(person, x) \wedge$ 
     $(\exists t_1, t_2) wantsCoffee(person, t_1, t_2) \wedge$ 
     $\neg status(person, Out)$  do
     $\pi(office : r)$ 
     $(optimize(\pi(person : office)$ 
       $(\exists t_1, t_2) wantsCoffee(person, t_1, t_2) \wedge$ 
       $\neg status(person, Out))?)$ ;
     $\pi t. [deliverCoffee(person, t)]$ ;
     $\pi t. [(now(t))?; goto(MO, t)]$ ;  $leaveItems$ );
     $optimize(\pi(v, t). [(now(t))?; sense(Battery, v, t)]$ ;
    if  $v < 24$  then  $chargeBattery(t)$  else  $Nil$ ].
end

```

According with the semantics of $\pi(x : \tau)p$ given in Section 4, the first occurrence of this construct in the procedure $deliver(r)$ represents the nondeterministic choice between offices given as the argument to the procedure, the second occurrence of this construct represents the nondeterministic choice between people occupying each office. We use the guard $(\exists t_1, t_2) wantsCoffee(person, t_1, t_2) \wedge \neg status(person, Out)?)$ to consider only those people who

want coffee and who are not out of their offices (we assume that initially all people are in their offices, but people may leave and come back at any time, and we do not model this behavior). Note that we use the operator $optimize$ to limit the scope of how far the incremental interpreter has to search to find the first person p to serve. Given the program $deliver(r)$, the interpreter considers all offices in r as alternative choices, makes a choice, chooses nondeterministically a person in that office and if the choice satisfies the guard, computes the optimal policy of delivering coffee to a chosen person. Because this last computation occurs inside the scope of the operator $optimize$, the interpreter does not look beyond. According to the reward function mentioned above and to the given probabilities that people are in their offices, the interpreter will decide who is the optimal person to serve first (it is the person with an early unfulfilled request and who is in her office with a high enough probability). Once this decision has been made, the robot can start executing the procedure $deliverCoffee(p, t)$ at some moment of time t (it will be determined as the side effect of computing the reward for giving coffee to p). Then, the procedure specifies to sense the voltage of the robot's battery: if the voltage is low, then the battery needs to be charged; otherwise, the computation continues until all requests are served.

```

proc deliverCoffee( $p, t$ )
  if  $robotLoc(MO)$  then
    if  $carrying(C)$  then  $serveCoffee(p, t)$ 
    else  $pickup(C, t)$ ;  $serveCoffee(p, t)$ 
    else if  $(\neg carrying(C))$  then  $goto(MO, t)$ ;
     $\pi t_1 [(now(t_1))?; pickup(C, t_1); serveCoffee(p, t_1)]$ 
    else  $serveCoffee(p, t)$ 
  end
proc serveCoffee( $p, t$ )
  if  $robotLoc(office(p))$  then  $give(C, p, t)$ 
  else  $goto(p, t)$ ;  $\pi t_1 [(now(t_1))?; give(C, p, t_1)]$ 
end
proc leaveItems
  while  $(\exists item) carrying(item)$  do
     $\pi item. [(carrying(item))?;$ 
     $\pi t_N (now(t_N))?; putBack(item, t_N)]$ .
end

```

The procedure $goto(l_2, t)$ commands the robot go from the current position l_1 determined by the fluent $robotLoc(l_1)$ to location l_2 beginning at time t . The procedure $goBetween$ commands the robot to go from the location l_1 to l_2 beginning at time t and taking Δ time units; it includes the sequence of $startGo$ and $endGo$.

```

proc goto( $l_2, t$ )  $(\pi l_1) [(robotLoc(l_1))$ ;
   $goBetween(l_1, l_2, travTime(l_1, l_2), t)]$ 
end
proc goBetween( $l_1, l_2, \Delta, t$ )
  if  $(l_1 = l_2 \wedge \Delta = 0)$  then  $Nil$ 
  else  $startGo(l_1, l_2, t)$ ;  $endGo(l_1, l_2, t + \Delta)$ 
end

```

To compare the computation time, we ran both the procedure $deliver(r)$ and its modification that does not include the construct $optimize$ (all runs were done on a computer with two Pentium II 300 Mz processors and 128Mb of RAM). More search is required if $deliver$ does not include the $optimize$ construct: it takes the incremental interpreter about 1 sec to compute the optimal policy for 2 people, about 12 sec to compute the optimal policy for 3 people and about 2 min to compute the optimal policy for 4 people (note that the decision tree grows exponentially with the number of people). But

when we run the procedure *deliver* with the *optimize* construct, then the computation time remains less than 3 sec even if the argument of *deliver* is the list of 4 or 5 offices with 7 or 8 people in them.

Notice that these programs restrict the policy that the robot can implement, leaving only one choice (the choice of person to whom to deliver an item) open to the robot and the rest of the robot's behavior is fixed by the program. Furthermore, these programs are, arguably, quite natural. Structuring a program this way may, in general, preclude optimal behavior. For instance, the outer most *optimize* construct restricts the search to find 'the best' person to serve at the current time without looking ahead what consequences this choice will have in the future. In circumstances when the robot is late this may lead the program to skip a next request in favor of a more later request. But if the program inside the scopes of the outer most *optimize* construct would consider all sequences of the next two requests (or the next three requests), then the choice of 'the best' person to serve now would take into account also a next one (or next two, respectively) request(s) to be served in the future and still the computation of the optimal policy will require an acceptable time. In addition, the program can be improved by allowing the robot to carry more than one item at a time: this may lead to a better behavior (for instance, in cases when two people from the same office requested different items). We will not delve into details here, because our intention was only to demonstrate the basic control structures of Golog programs, a new construct *optimize* and to give the feeling of how much time is required to compute optimal policies.

6 Discussion

The incremental Golog interpreter based on the single-step *Trans*-semantics is introduced in [De Giacomo and Levesque, 1999b]. The Golog programs considered there may include binary sensing actions. The interpreter considered in our paper is motivated by similar intuitions, but it is based on a different decision-theoretic semantics and employs more expressive representation of sensing. The paper [De Giacomo and Levesque, 1999a] introduces guarded sensed fluent axioms (GSFA) and guarded successor state axioms (GSSA) and assumes that there is a stream of sensor readings available at any time. These readings are represented by unary sensing functions (syntactically, they look like functional fluents). Because we introduce the representation for sensing actions, they can be mentioned explicitly in Golog programs or can be executed by the interpreter. The major advantage of our representation is that it does not require consistency of sensory readings with the action theory (this may prove useful in solving diagnostic tasks: [McIlraith, 1998]). The execution monitoring framework proposed in [De Giacomo *et al.*, 1998] assumes that the feedback from the environment is provided in terms of actions executed by other agents. Because in this paper we assume that the feedback is provided in terms of sensory readings, this may lead to development of a more realistic framework. An approach to integrating planning and execution in stochastic domains [Dearden and Boutilier, 1994] is an alternative to the approach proposed here.

7 Concluding remarks

Several important issues are not covered in this paper. One of them is monitoring and rescheduling of policies. Note that all actions in policies have time arguments which will be in-

stantiated by moments of time: when the incremental interpreter computes an optimal policy, it also determines a schedule when actions have to be executed. But in realistic scenarios, when the robot is involved in ongoing processes extended over time, it may happen that a process will terminate earlier or later than it was expected.

The diagnostic task that the current version of the on-line interpreter solves is admittedly oversimplified. We expect that additional research integrating the on-line incremental interpreter with the approach proposed in [McIlraith, 1998] will allow us to formulate a more comprehensive version.

Successful tests of the implementation described here were conducted on the mobile robot in a real office environment (joint work with Sam Kaufman). They demonstrated that using the expressive set of Golog operators it is straightforward to encode domain knowledge as constraints on the given large MDP problem. The operator *optimize(p)* proved to be useful in providing heuristics which allowed to compute sub-policies in real time.

References

- [Bacchus *et al.*, 1995] Fahiem Bacchus, Joseph Y. Halpern, and Hector J. Levesque. Reasoning about noisy sensors in the situation calculus. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1933–1940, Montreal, 1995.
- [Boutilier *et al.*, 2000] C. Boutilier, R. Reiter, M. Soutchanski, and S. Thrun. Decision-theoretic, high-level robot programming in the situation calculus. In *Proc. of the 17th National Conference on Artificial Intelligence (AAAI'00)*, Austin, Texas, 2000. Available at: <http://www.cs.toronto.edu/~cogrobo/>.
- [De Giacomo and Levesque, 1999a] G. De Giacomo and H. Levesque. Projection using regression and sensors. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, 1999.
- [De Giacomo and Levesque, 1999b] G. De Giacomo and H.J. Levesque. An incremental interpreter for high-level programs with sensing. In Levesque and Pirri, editors, *Logical Foundations for Cognitive Agents: Contributions in Honor of Ray Reiter*, pages 86–102. Springer, 1999.
- [De Giacomo *et al.*, 1998] G. De Giacomo, R. Reiter, and M.E. Soutchanski. Execution monitoring of high-level robot programs. In *Principles of Knowledge Representation and Reasoning: Proc. of the 6th International Conference (KR'98)*, pages 453–464, Trento, Italy, 1998.
- [Dearden and Boutilier, 1994] Richard Dearden and Craig Boutilier. Integrating planning and execution in stochastic domains. In *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, pages 162–169, 1994.
- [Funge, 1998] J. Funge. *Making Them Behave: Cognitive Models for Computer Animation*, Ph.D. Thesis. Dept. of Computer Science, Univ. of Toronto, 1998.
- [Grosskreutz, 2000] H. Grosskreutz. Probabilistic projection and belief update in the pgolog framework. In *The 2nd International Cognitive Robotics Workshop, 14th European Conference on AI*, pages 34–41, Berlin, Germany, 2000.
- [Lakemeyer, 1999] G. Lakemeyer. On sensing and off-line interpreting in golog. In Levesque and Pirri, editors, *Logical Foundations for Cognitive Agents: Contributions in Honor of Ray Reiter*, pages 173–189. Springer, 1999.
- [Levesque, 1996] H.J. Levesque. What is planning in the presence of sensing? In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, volume 2, pages 1139–1145, Portland, Oregon, 1996.

- [McIlraith, 1998] S. McIlraith. Explanatory diagnosis: Conjecturing actions to explain observations. In *Principles of Knowledge Representation and Reasoning: Proc. of the 6th International Conference (KR'98)*, pages 167–177, Italy, 1998.
- [Pirri and Finzi, 1999] F. Pirri and A. Finzi. An approach to perception in theory of actions: part 1. *Linköping Electronic Articles in Computer and Information Science*, 4(41), 1999.
- [Pirri and Reiter, 1999] F. Pirri and R. Reiter. Some contributions to the metatheory of the situation calculus. *Journal of the ACM*, 46(3):261–325, 1999.
- [Reiter, 1998] R. Reiter. Sequential, temporal golog. In *Principles of Knowledge Representation and Reasoning: Proc. of the 6th International Conference (KR'98)*, pages 547–556, Trento, Italy, 1998. Available at: <http://www.cs.toronto.edu/~cogrobo/>.
- [Reiter, 2000] R. Reiter. *Knowledge in Action: Logical Foundations for Describing and Implementing Dynamical Systems*. <http://www.cs.toronto.edu/~cogrobo/>, 2000.
- [Scherl and Levesque, 1993] R. Scherl and H.J. Levesque. The frame problem and knowledge producing actions. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 689–695, Washington, DC, 1993.