
Explanatory Diagnosis: Conjecturing actions to explain observations

Sheila A. McIlraith*

Knowledge Systems Lab
Stanford University
Stanford, CA 94305-9020
sam@ksl.stanford.edu

Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94301
mcilraith@parc.xerox.com

Abstract

In this paper we present contributions towards a logical theory of diagnosis for systems that can be affected by the actions of agents. Specifically, we examine the task of conjecturing diagnoses to explain *what happened* to a system, given a theory of system behaviour and some observed (aberrant) behaviour. We characterize what happened by introducing the notion of explanatory diagnosis in the language of the situation calculus. Explanatory diagnoses conjecture sequences of actions to account for a change in system behaviour. As such, we show that determining an explanatory diagnosis is analogous to classical AI planning with state constraints and incomplete knowledge. The representation scheme we employ provides an axiomatic solution to the frame, ramification and qualification problems. Exploiting this representation, we show that determining an explanatory diagnosis can be achieved by regression followed by theorem proving in the database describing what is known of the initial state of our system. Further, we show that by exploiting features inherent to diagnosis problems, we can simplify the diagnosis task.

1 INTRODUCTION

Given a theory of system behaviour and some observed aberrant behaviour, the traditional objective of diagnosis is to conjecture *what is wrong* with the system, (e.g., which components of the device are behaving abnormally, what diseases the patient is suffering from,

etc.). Each candidate diagnosis consists of a subset of distinguished literals that are conjectured to be true or false in order to account for the observation in some way. Different criteria have been proposed for determining the space of such candidate diagnoses. Within formal accounts of diagnosis, two widely accepted definitions of diagnosis are consistency-based diagnosis (e.g., (Reiter 1987), (de Kleer et al. 1992)), and abductive explanation (e.g., (de Kleer et al. 1992), (Poole 1988), (Console and Torasso 1991), (McIlraith 1994)). Such research has historically focussed on static systems. Recently, some researchers have advocated extending diagnostic problem solving (DPS) to enable reasoning about actions, under the argument that DPS is purposive in nature and that systems operate within and are affected by agents¹.

In this paper we focus upon one aspect of diagnosing such dynamic systems. In particular, given a theory of system behaviour and some observation of (aberrant) behaviour, our concern is with the task of conjecturing diagnoses to explain *what happened* to the system (i.e., what actions or events occurred to result in the observed behaviour) (e.g., (Cordier and Thiébaux 1994), (McIlraith 1994a)). Knowing or conjecturing what happened is interesting in its own right, but it can also help to further constrain the space of possible states of the system. In so doing, conjecturing what happened assists in the prediction of what is wrong with a system, as well as predicting other relevant system behaviour. Compared to our traditional notion of *what is wrong* diagnoses, knowing *what happened* can more accurately capture the root cause of system malfunction rather than its manifestations, thus providing for the identification of future preventative as well as prescriptive actions.

In the spirit of previous foundational work in model-

* Much of this research was performed while the author was a doctoral candidate at the University of Toronto, Canada.

¹An agent could be another system, a robot, a human, or nature.

based diagnosis (MBD) (e.g., (Reiter 1987), (Console and Torasso 1991), (de Kleer et al. 1992)), this paper presents a logical characterization of the diagnosis task. We take as our starting point the existing MBD research on characterizing diagnoses for static systems *without* a representation of actions (e.g., (de Kleer et al. 1992), (Console and Torasso 1991), (Reiter 1987)). Next, we exploit a situation calculus representation scheme previously proposed by the author (McIlraith 1997a) that enables the integration of a representation of action with the representation of the behaviour of a static system. With this representation in hand, we provide a logical characterization for the task of determining *what happened* to a system. The characterization is presented in the guise of *explanatory diagnosis*.

The distinguishing features of our characterization are afforded in great part by the richness of our representation scheme which provides a comprehensive and semantically justified representation of action and change. In particular, our representation provides an axiomatic closed-form solution to the frame and ramification problems, thus capturing the direct and indirect effects of actions in a compiled representation. This is critical to the ability to generate explanatory diagnoses efficiently. Further, our representation provides a closed-form solution to the qualification problem, thus identifying the conditions under which an action is possible. It is interesting to note that when we are dealing with incomplete knowledge of our initial state, conjecturing an action or sequence of actions also requires conjecturing that its preconditions are satisfied, which in many instances serves to further constrain our search.

As we show in the sections to follow, our characterization establishes a direct link between explanatory diagnosis and planning, deductive plan synthesis, and abductive planning. As a consequence of a completeness assumption embedded in our representation, we can exploit goal-directed reasoning in the form of regression (Waldinger 1977) in order to generate diagnoses. This completeness assumption also provides for an easy mapping of our situation calculus representation to Prolog. While explanatory diagnoses can be mathematically characterized in an analogous fashion to plans, an important distinction of explanatory diagnoses is that they can be refined to exploit features of our diagnosis problem that have no meaning in the context of planning. Indeed, we use diagnosis-specific attributes to define variants of explanatory diagnosis to deal with the challenges of incomplete knowledge and large search spaces.

2 REPRESENTATION SCHEME

2.1 SITUATION CALCULUS LANGUAGE

The situation calculus language in which we axiomatize our domains is a sorted first-order language with equality. The sorts are of type \mathcal{A} for primitive *actions*, \mathcal{S} for *situations*, and \mathcal{D} for everything else, including domain objects (Lin and Reiter 1994). We represent each action as a (possibly parameterized) first-class object within the language. Situations are simply sequences of actions. The evolution of the world can be viewed as a tree rooted at the distinguished initial situation S_0 . The branches of the tree are determined by the possible future situations that could arise from the realization of particular sequences of actions. As such, each situation along the tree is simply a history of the sequence of actions performed to reach it. The function symbol *do* maps an action term and a situation term into a new situation term. For example, $do(turn_on_pmp, S_0)$ is the situation resulting from performing the action of turning on the pump in situation S_0 . The distinguished predicate $Poss(a, s)$ denotes that an action a is possible to perform in situation s (e.g., $Poss(turn_on_pmp, S_0)$). Thus, $Poss$ determines the subset of the situation tree consisting of situations that are possible in the world. Finally, those properties or relations whose truth value can change from situation to situation are referred to as *fluents*. For example, the fluent $on(Pmp, s)$ expresses that the pump is on in situation s .

The dialect of the situation calculus that we use in this paper is restricted to primitive, determinate actions. Our language does not include functional fluents, nor does it include a representation of time, concurrency, or complex actions, but we believe the results presented herein can be extended to more expressive dialects of the situation calculus (e.g., (Reiter 1996)).

2.2 DOMAIN REPRESENTATION

In this section we overview the representation scheme we use to characterize the system we will be diagnosing. The scheme, proposed in (McIlraith 1997a), integrates a situation calculus theory of action with a MBD system description, SD (de Kleer et al. 1992). The resulting representation of a system comprises both domain-independent and domain-specific axioms. The domain-independent axioms are the foundational axioms of the discrete situation calculus, Σ_{found} (Lin and Reiter 1994). They are analogous to the axioms of Peano arithmetic, modified to define the branching structure of our situation tree, rather than the number line. The domain-specific axioms, T specify both the

behaviour of the static system, and the actions² that can affect the state of the system, as well as those actions required to achieve testing and repair. Together they define our situation calculus representation $\Sigma = \Sigma_{found} \wedge T$.

The domain-specific axioms composing T result from application of a procedure proposed in (McIlraith 1997a) that compiles a typical MBD system description, SD and a set of axioms relating to the preconditions and effects of actions into a representation that provides a closed-form solution to the frame, ramification and qualification problems. The resulting domain axiomatization $T = T_{SC}^{S_0} \wedge T_{domain} \wedge T_{SS} \wedge T_{AP} \wedge T_{UNA} \wedge T_{DCA} \wedge T_{S_0}$ is described below. The representation is predicated on an explicit causal ordering of fluents and an a completeness assumption. The assumption states that all the conditions underwhich an action a can lead, directly or indirectly, to fluent F becoming true or false in the successor state are captured in the axiomatization of our system.

We illustrate the representation scheme in terms of a small portion of a power plant feedwater system (McIlraith 1997) derived from the APACS project (Kramer et al. 1996). This simplified example models the filling of a vessel either by the operation of an electrically powered (Pwr) pump (Pmp), by manual filling, or by a siphon that was started by the pump or by manual filling. For notational convenience, all formulae are understood to be universally quantified with respect to their free variables, unless explicitly indicated otherwise. For a more thorough description of this representation scheme, and for a more extensive example please see ((McIlraith 1997a) and (McIlraith 1997)).

Every domain axiomatization, T comprises the following sets of axioms.

$$T_{SC}^{S_0} \wedge T_{domain} \wedge T_{SS} \wedge T_{AP} \wedge T_{UNA} \wedge T_{DCA} \wedge T_{S_0}$$

The set of state constraints relativized to situation S_0 , $T_{SC}^{S_0}$ capture what is implicitly true about the initial database. They can be acquired from a typical MBD system description, SD as described in (McIlraith 1997a). In our simple example, $T_{SC}^{S_0}$ is as follows

$$\begin{aligned} \neg AB(Pwr, S_0) \wedge \neg AB(Pmp, S_0) \wedge on(Pmp, S_0) \\ \supset filling(S_0) \quad (1) \\ manual_fill(S_0) \supset filling(S_0) \quad (2) \\ \neg(on(Pmp, S_0) \wedge manual_fill(S_0)) \quad (3) \end{aligned}$$

The set of domain constraints, T_{domain} is as follows.

$$Pwr \neq Pmp \quad (4)$$

²Actions can be performed by agents: a human, another system, or nature.

The set of successor state axioms, T_{SS} is composed of axioms of the following general form, one for each fluent F .

$$Poss(a, s) \supset [F(do(a, s)) \equiv \Phi_F] \quad (5)$$

where Φ_F is a simple formula³ of a particular syntactic form. Intuitively, a successor state axiom says the following:

$$\begin{aligned} Poss(a, s) \supset [fluent(do(a, s)) \equiv \\ \text{an action made it true} \\ \vee \text{ a state constraint made it true} \\ \vee \text{ it was already true} \\ \wedge \text{ neither an action nor a state constraint} \\ \text{made it false}] \end{aligned}$$

The following axioms compose T_{SS} for our example.

$$\begin{aligned} Poss(a, s) \supset [on(Pmp, do(a, s)) \equiv a = turn_on_pmp \\ \vee (on(Pmp, s) \wedge a \neq turn_off_pmp)] \quad (6) \end{aligned}$$

$$\begin{aligned} Poss(a, s) \supset [AB(Pwr, do(a, s)) \equiv a = pwr_failure \\ \vee (AB(Pwr, s) \wedge a \neq aux_pwr \wedge a \neq pwr_fix)] \quad (7) \end{aligned}$$

$$\begin{aligned} Poss(a, s) \supset [AB(Pmp, do(a, s)) \equiv a = pmp_burn_out \\ \vee (AB(Pmp, s) \wedge a \neq pmp_fix)] \quad (8) \end{aligned}$$

$$\begin{aligned} Poss(a, s) \supset [manual_fill(do(a, s)) \equiv \\ a = turn_on_manual_fill \\ \vee (manual_fill(s) \\ \wedge a \neq turn_off_manual_fill)] \quad (9) \end{aligned}$$

$$\begin{aligned} Poss(a, s) \supset [filling(do(a, s)) \equiv \\ a = turn_on_manual_fill \\ \vee (manual_fill(s) \wedge a \neq turn_off_manual_fill) \\ \vee [(a \neq pwr_failure \\ \wedge (\neg AB(Pwr, s) \vee a = aux_pwr \\ \vee a = pwr_fix)) \\ \wedge (a \neq pmp_burn_out \\ \wedge (\neg AB(Pmp, s) \vee a = pmp_fix)) \\ \wedge (a = turn_on_pmp \\ \vee (on(Pmp, s) \wedge a \neq turn_off_pmp))] \\ \vee (filling(s) \wedge a \neq stop_siphon)] \quad (10) \end{aligned}$$

Axiom (6) states that if action a is possible in situation s , then the pump is on in the situation resulting from performing action a in situation s

³A simple formula only mentions domain-specific predicate symbols, fluents do not include the function symbol do , there is no quantification over sort *situation*, and there is at most one free *situation* variable.

(i.e., $on(Pmp, do(a, s))$) if and only if the action a is $turn_on_pmp$, or the pump was already on in s and a was not the action $turn_off_pmp$.

The set of action precondition axioms, T_{AP} is composed of axioms of the following general form, one for each action prototype A in the domain.

$$Poss(A(\vec{x}), s) \equiv \Pi_A \quad (11)$$

where Π_A is a simple formula with respect to s , capturing the necessary and sufficient conditions for prototype action $A(\vec{x})$ to be executable in situation s .

$$Poss(stop_siphon, s) \equiv (\neg manual_fill(s) \wedge \neg on(Pmp, s)) \quad (12)$$

$$Poss(pmp_fix, s) \equiv \neg on(Pmp, s) \quad (13)$$

$$Poss(pmp_burn_out, s) \equiv on(Pmp, s) \quad (14)$$

$$Poss(turn_on_manual_fill, s) \equiv \neg on(Pmp, s) \quad (15)$$

$$Poss(turn_on_pmp, s) \equiv \neg manual_fill(s) \quad (16)$$

$$Poss(turn_off_pmp, s) \equiv on(Pmp, s) \quad (17)$$

$$Poss(turn_off_manual_fill, s) \equiv manual_fill(s) \quad (18)$$

$$Poss(pwr_failure, s) \equiv \quad (19)$$

$$Poss(pwr_fix, s) \equiv Poss(aux_pwr, s) \equiv \text{true} \quad (20)$$

Finally, we provide a possible set of initial conditions for our system, T_{S_0} . These constitute the explicit aspect of the initial database. Note that in general we do not have complete knowledge of the initial state of our system. This makes the task of diagnosis all the more challenging. In this example, we do not know initially whether the pump and power are operating normally. We also do not know whether the vessel was filling in the initial state.

$$on(Pmp, S_0) \wedge \neg manual_fill(S_0) \quad (21)$$

T_{UNA} and T_{DCA} are the unique names axioms for actions and the domain closure axiom for actions, respectively.

This concludes the description of our representation scheme. Before advancing to issues of diagnosis, we note (Reiter 1998) that our proposed situation calculus representation can be viewed as an executable specification because it is easily realized in Prolog by exploiting Prolog's completion semantics and simply replacing the equivalence connectives characteristic of axioms in T_{SS} and T_{AP} by implication connectives. The Lloyd-Topor transformation (Lloyd 1987) must then be applied to convert this theory into Prolog clausal form. Later in this paper, we will advocate using Waldinger's notion of regression to rewrite axioms of our representation and simplify computation. This type of regression rewriting is precisely achieved by Prolog's backwards chaining mechanism.

3 PRELIMINARIES

With our representation in hand, we turn our attention to the task of diagnosis. In this section we introduce the framework for performing diagnosis relative to our representation. For our purposes we adopt the ontological and notational convention of the MBD literature and view the systems we are diagnosing as comprising a number of interacting *components*, *COMPS*. These components have the property of being either abnormal or normal in a situation. We express this property in our situation calculus language using the fluent AB . For example, $AB(Pmp, s)$ denotes that the pump component is abnormal in situation s . Note that the use of AB is not mandatory to the contributions of this paper. Once again, following the convention in the MBD literature, we define our diagnoses relative to the domain-independent concept of a *system* (de Kleer et al. 1992), adapted to our situation calculus framework.

Definition 1 (System)

A system is a quadruple $(\Sigma, HIST, COMPS, OBS)$ where:

- Σ , the background theory, is a set of situation calculus sentences describing the behaviour of our system and the actions that can affect it.
- $HIST$, the history, is a sequence of ground actions $[a_1, \dots, a_k]$ that were performed starting in S_0 .
- $COMPS$, the components, is a finite set of constants.
- OBS_F , the observation, is a simple formula composed of fluents whose only free variable is the situation variable s , and which are otherwise ground.

Example 1

In our power plant example above, Σ is our axiomatization $\Sigma_{found} \wedge T$ and $COMPS = \{Pmp, Pwr\}$. The observation, OBS_F could be $filling(s)$, for example. $HIST$ could be empty, i.e., $[\]$, or perhaps $[turn_on_pmp]$.

4 EXPLANATORY DIAGNOSIS

In this section we introduce and formally characterize the notion of an explanatory diagnosis which conjectures *what happened* to result in some observed (aberrant) behaviour. Given a system, $(\Sigma, HIST, COMPS, OBS_F)$, the objective of explanatory diagnosis is to conjecture a sequence of actions, $[\alpha_1, \dots, \alpha_n]$

such that our observation is true in the situation resulting from performing that sequence of actions in $do(HIST, S_0)$.

Definition 2 (Explanatory Diagnosis)

An explanatory diagnosis for system $(\Sigma, HIST, COMPS, OBS_F)$ is a sequence of actions $E = [\alpha_1, \dots, \alpha_n]$ such that,

- $\Sigma \models Poss(HIST \cdot E, S_0)^3$
 $\wedge OBS_F(do(HIST \cdot E, S_0))$.

Thus, E is an explanatory diagnosis when the observation is true in the situation resulting from performing the sequence of actions E in situation $do(HIST, S_0)$, and further that the preconditions for each action of the action sequence $HIST \cdot E$ are true in the appropriate situations, commencing at S_0 .

The problem of determining explanatory diagnoses is an instance of temporal explanation or postdiction (e.g., (Shanahan 1993)), and is related to the classical AI planning problem. In particular, identifying the sequence of actions composing an explanatory diagnosis, E is analogous to the plan synthesis problem, and thus is realizable using deduction on the situation calculus axioms. According to (Green 1969), a plan to achieve a goal $G(s)$ is obtained as a side effect of proving $Axioms \models \exists s.G(s)$. The binding for the situation variable s represent the sequence of actions. In our case, $Axioms \models \exists s.G(s)$ is analogous to $\Sigma \models \exists s.OBS_F(s)$. As such, our representation enables us to generate explanatory diagnoses deductively, just as we could deductively generate a plan in the situation calculus. Note that the task of generating explanatory diagnoses is analogous to plan synthesis in the presence of state constraints – a challenging problem. Our representation scheme eliminates the additional challenges presented by state constraints by providing a domain axiomatization that a priori solves the frame, ramification and qualification problems.

Example 2

Continuing with our power plant example, given the system $(\Sigma, [], \{Pwr, Pmp\}, \neg filling(s))$, the sequence of actions $[pwr_failure]$ constitutes one example of an explanatory diagnoses for the system. Another explanatory diagnosis for our system is $[turn_off_pmp]$.

³**Notation:**

$HIST \cdot E$ is an abbreviation for $[a_1, \dots, a_k, \alpha_1, \dots, \alpha_n]$.
 $do([a_1, \dots, a_m], s)$ is an abbreviation for
 $do(a_m, (do(a_{m-1}, (do(a_{m-2}, (\dots, (do(a_1, s))))))))$.
Finally, $Poss([a_1, \dots, a_n], s)$ is an abbreviation
for $Poss(a_1, s) \wedge Poss(a_2, do(a_1, s)) \wedge \dots$
 $\wedge Poss(a_n, do([a_1, \dots, a_{n-1}], s))$.

Observe that for certain problems there can be an infinite number of sequences of actions that constitute explanatory diagnoses. For example, the following sequences of actions also constitute valid explanatory diagnoses for our example system:

$[pwr_failure, pwr_fix, pwr_failure]$,
 $[pwr_failure, aux_pwr, pwr_failure]$,
 $[turn_off_pmp, pwr_failure, turn_on_pmp]$,

and so on.

Definition 2 is not sufficiently discriminating to eliminate these, clearly suboptimal explanatory diagnoses. We must define a preference criterion. Probability measures, even simple order of magnitude probabilities have provided an effective preference criterion for many applications of MBD (de Kleer 1991). Likewise, in the case of determining explanatory diagnoses in the context of the situation calculus, probabilities will serve us well in identifying preferred explanatory diagnoses. Unfortunately, probability measures are not always available and the correct treatment of probabilities in our situation calculus framework is only now being developed. In this paper, we limit our discussion to what we refer to as a chronologically simple preference criterion.

In our chronologically simple preference criterion, we prefer diagnoses that are relativized to situations reached without performing any extraneous actions. Note that this preference criterion is syntactic in nature, relying on the notion of a primitive action as a unit measure.

Definition 3 (Simpler)

Given a sequence of actions $HIST = [\alpha_1, \dots, \alpha_n]$, define $ACTS(HIST)$ to be the set $\{\alpha_1, \dots, \alpha_n\}$, and $LEN(HIST)$ to be the length of the sequence of actions composing $HIST$.

Thus, given $HISTA = [a_1, \dots, a_n]$ and $HISTB = [b_1, \dots, b_n]$, situation $S_A = do(HISTA, S_0)$ is simpler than situation $S_B = do(HISTB, S_0)$ iff $ACTS(HISTA) \subseteq ACTS(HISTB)$ and $LEN(HISTA) < LEN(HISTB)$.

Definition 4 (Chronologically Simple Diag.)

E is a chronologically simple explanatory diagnosis for system $(\Sigma, HIST, COMPS, OBS_F)$ iff E is an explanatory diagnosis for the system, and there is no explanatory diagnosis E' such that situation $S' = do(HIST \cdot E', S_0)$ is simpler than situation $S = do(HIST \cdot E, S_0)$.

We might further distinguish this criterion to prefer chronologically simple explanatory diagnoses comprised solely of actions performed by nature.

Finally, observe that the characterization of explanatory diagnosis just presented assumes that E and OBS_F occur *after HIST*. While this assumption is not critical to characterizing explanatory diagnoses, it acts as a form of preference, facilitating computation of E .

5 EXPLOITING REGRESSION

In the previous section, we provided a characterization of explanatory diagnosis. We observed that computing an explanatory diagnosis for system $(\Sigma, HIST, COMPS, OBS_F)$ is analogous to generating a plan to achieve a goal $OBS_F(s)$ starting with axioms Σ and situation $do(HIST, S_0)$. At first glance, the general problem of computing explanatory diagnoses does not look very promising for at least three reasons: the second-order induction axiom in Σ_{found} , the potential incompleteness of the initial database, and the potentially large size of the situation search space. In this section, we show how diagnoses can be computed by exploiting regression (Waldinger 1977). We are able to exploit regression as a direct result of the embedded completeness assumption inherent in our representation scheme. In this context, regression is a recursive rewriting procedure that we use to reduce the nesting of the *do* function in situation terms, or to eliminate the *Poss* predicate. We show that generating explanatory diagnoses reduces to regression followed by entailment with respect to the initial database. Computationally, the merit of regression is that it searches backwards through the situation space from the observation rather than searching forward from the initial database. Under the assumption that the observation consists of fewer literals than the initial database, regression will make for more efficient search. Observe that Prolog's backwards chaining mechanism achieves the substitution performed by regression.

Following directly in the spirit of previous work by Reiter ((Reiter 1991), (Reiter 1992)) and more recently (Reiter 1998) on the exploitation of regression for planning and query answering, we first define two regression operators, \mathcal{R}^* and \mathcal{R}_{Poss} .

Definition 5 (Regression Operator \mathcal{R}^*) *Given a set of successor state axioms, T_{SS} composed of axioms of the form $Poss(a, s) \supset [F(do(a, s)) \equiv \Phi_F]$, $\mathcal{R}^*[\Psi]$, the repeated regression of formula Ψ with respect to successor state axioms T_{SS} is the formula that is obtained from Ψ by repeatedly replacing each fluent $F(do(a, s))$ in Ψ by Φ_F , until the resulting formula makes no mention of the function symbol *do*.*

For example,

$$\begin{aligned} & \mathcal{R}^*[on(Pmp, do(turn_on_manual_fill, do(turn_on_pmp, S_0)))] \\ &= \mathcal{R}^*[(turn_on_manual_fill = turn_on_pmp) \\ & \vee (on(Pmp, do(turn_on_pmp, S_0)) \\ & \wedge (turn_on_manual_fill \neq turn_off_pmp))] \\ &= \mathcal{R}^*[false \vee (on, (Pmp, do(turn_on_pmp, S_0)) \wedge true)] \\ &= \mathcal{R}^*[(turn_on_pmp = turn_on_pmp) \\ & \vee (on(pmp, S_0)) \wedge (turn_on_pmp \neq turn_off_pmp))] \\ &= true \end{aligned}$$

We can similarly define a *Poss* regression operator over the set of action precondition axioms, T_{AP} . This regression operation rewrites each occurrence of the literal $Poss(a, s)$ by Π_A as defined in the action precondition axioms.

Definition 6 (Regression Operator \mathcal{R}_{Poss})

Given a set of action precondition axioms, T_{AP} composed of axioms of the form $Poss(A(\vec{x}), s) \equiv \Pi_A$, $\mathcal{R}_{Poss}[W]$ is the formula obtained by replacing each occurrence of predicate $Poss(A(\vec{x}), s)$ by Π_A . All other literals of W remain the same.

Reiter proved soundness and completeness results for regression applied to a theory with no state constraints (Theorem 1, Theorem 2, (Reiter 1992)). In the theorem below, we prove soundness and completeness results for our representation scheme which includes state constraints. The theory Σ_{init} mentioned in the theorem below is a subset of Σ containing only information about the initial situation, and no information about successor situations. It also excludes the induction axiom of Σ_{found} .

Theorem 1 (Soundness & Completeness)

Given

- Σ_{init} , a subset of the situation calculus theory Σ , such that $\Sigma_{init} = \Sigma_{UNS} \wedge T_{S_0} \wedge T_{SC}^{S_0} \wedge T_{domain} \wedge T_{UNA}$, where Σ_{UNS} is a subset of Σ_{found} containing the set of unique names axioms for situations.
- a sequence of ground actions, s_HIST such that $\Sigma_{init} \wedge \mathcal{R}^*[\mathcal{R}_{Poss}[Poss(s_HIST, S_0)]]$ is satisfiable.
- $Q(s)$, a simple formula whose only free variable is the situation variable s .

Suppose $S = do(s_HIST, S_0)$, then

- $\Sigma \models Q(do(s_HIST, S_0))$ iff $\Sigma_{init} \models \mathcal{R}^*[Q(do(s_HIST, S_0))]$,
- $\Sigma \models Poss(s_HIST, S_0)$ iff $\Sigma_{init} \models \mathcal{R}^*[\mathcal{R}_{Poss}[Poss(s_HIST, S_0)]]$,

- $\Sigma \wedge Poss(s_HIST, S_0) \wedge Q(do(s_HIST, S_0))$ is satisfiable iff $\Sigma_{init} \wedge \mathcal{R}^*[\mathcal{R}_{Poss}[Poss(s_HIST, S_0)]] \wedge \mathcal{R}^*[Q(do(s_HIST, S_0))]$ is satisfiable.

Thus, assuming situation s is a possible situation and exploiting regression, $Q(s)$ holds at situation s iff its regression is entailed in the initial database. The beauty of Theorem 1 is that it enables us to generate explanatory diagnoses via regression followed by theorem proving in the initial database, without the need for the second-order induction axiom in Σ_{found} . From these results, we can characterize explanatory diagnosis with respect to regression.

Corollary 1 (Expl. Diagnosis w/ Regression)

The sequence of actions $E = [\alpha_1, \dots, \alpha_k]$ is an explanatory diagnosis for system $(\Sigma, HIST, COMPS, OBS_F)$ iff

- $\Sigma_{init} \models \mathcal{R}^*[\mathcal{R}_{Poss}[Poss(HIST \cdot E, S_0)]] \wedge \mathcal{R}^*[OBS_F(do(HIST \cdot E, S_0))]$.

6 EXPLOITING THE TASK

In the previous section we showed that we could exploit regression to simplify the computation of explanatory diagnoses. A remaining source of difficulty in generating explanatory diagnoses is that our search space may be large and our initial database may be incomplete. An incomplete initial database both under-constrains our search problem, and precludes us from using certain planning machinery, such as STRIPS (Fikes and Nilsson 1971), that assumes a complete or near-complete initial database (Lin and Reiter 1995). In this section we show how to exploit features of diagnosis problems to further assist in the generation of explanatory diagnoses. In particular, we propose to 1) make assumptions regarding our domain, 2) relax our criteria for explanatory diagnoses, and 3) verify rather than generate diagnoses. An additional means of simplifying our computational task is to use likelihoods of actions and action sequences to focus search for explanatory diagnoses. Detailed discussion of this option is beyond the scope of this paper.

6.1 ASSUMPTION-BASED DIAGNOSES

In diagnostic problem solving it is common to make further assumptions that are consistent with what we know of the world. For example, we may assume that in the absence of information to the contrary, components are operating normally, or certain properties hold of the world. To support such assumption-based reasoning, we define the notion of an assumption-based explanatory diagnosis.

Definition 7 (Assumption-based Expl. Diag.)
Given an assumption $H(S)$ relativized to ground situation S such that

- $S_0 \leq^4 S \leq do(HIST \cdot E, S_0)$,
- $\Sigma \wedge H(S)$ is satisfiable, and
- $\Sigma \wedge H(S) \models Poss(HIST, S_0)$.

An assumption-based explanatory diagnosis for system $(\Sigma, HIST, COMPS, OBS_F)$ under assumption $H(S)$ is a sequence of actions $E = [\alpha_1, \dots, \alpha_k]$ such that,

- $\Sigma \wedge H(S) \models Poss(HIST \cdot E, S_0) \wedge OBS_F(do(HIST \cdot E, S_0))$.

In Example 2 of the previous section, we did not have complete information about the initial state of our system. It could actually have been the case that observation $\neg filling$ was true in S_0 , i.e., $\neg filling(S_0)$, but since it was not entailed by Σ , the empty action sequence was not proposed as a valid explanatory diagnosis, and we were forced to conjecture a sequence of actions to account for our observation. If we assume $\neg filling(S_0)$, then the empty sequence of actions is indeed an assumption-based explanatory diagnosis.

In generating explanatory diagnoses, we may want to make a priori assumptions about the world, conjoin these assumptions to our theory and then try to compute our explanatory diagnoses. For example, we may wish to assume that all components are operating normally in S_0 . This would be achieved by making $H(S)$ in our definition above equal to $\bigwedge_{c \in COMPS} \neg AB(c, S_0)$ (i.e., $\neg AB(Pmp, S_0) \wedge \neg AB(Pwr, S_0)$). Similarly, we may wish to assume that the observation, OBS_F is true in $do(HIST, S_0)$. In our example above, this would mean assuming $\neg filling(S_0)$. In other instances, we might want $H(S)$ to equal a *what is wrong* diagnosis that we are currently entertaining, relativized to a previous situation, e.g., $AB(Pmp, S_0)$.

In still other instances, we may not want to fix our assumptions a priori but rather make the minimum number of assumptions necessary to generate a chronologically simple explanatory diagnosis. Such assumptions might be limited to a distinguished set of literals which the domain axiomatizer considers to be legitimately assumable (e.g., AB fluents). There might also be a partial ordering on such assumables fluents.

⁴**Notation:** The transitive binary relation $<$ defined in Σ_{found} further limits our situation tree by restricting the actions that are applied to a situation to those whose preconditions are satisfied in the situation. Intuitively, if $s < s'$, then s and s' are on the same branch of the tree with s closer to S_0 than s' . Further, s' can be obtained from s by applying a sequence of actions whose preconditions are satisfied by the truth of the $Poss$ predicate.

The distinction regarding when we make our assumptions affects the machinery by which we compute assumption-based explanatory diagnoses. If assumptions are made prior to computation we simply conjoin the regression of the assumption to the initial database and use regression and theorem proving as we would for generating normal explanatory diagnoses. When assumptions are interleaved with computation, generating an assumption-based explanatory diagnosis requires abduction.

In a theorem prover, abduction is generally implemented as proof-tree completion, i.e., by resolving dead-ends of proof trees with abducible literals. To generate an explanatory diagnosis, an abductive theorem prover would attempt to prove $OBS_F(s)$. If an attempted proof failed because it dead-ended on a literal or literals that were assumable, then these would be abducted and the proof continued. We have not yet implemented such an abduction engine for our situation calculus system. It is interesting to note that in the context of planning, (Shanahan 1997) has used abduction to implement a partial order planner for the event calculus. In contrast to the deductive approach of the situation calculus, this planner abduces actions and the relative order of certain actions.

6.2 POTENTIAL DIAGNOSES

In addition to making assumptions to help complete our theory, we can also facilitate computation by relaxing the criteria for defining an explanatory diagnosis. To this end, we observe that the requirement in Definitions 2 and 7 that $\Sigma \models Poss(HIST \cdot E, S_0)$ may be too stringent in the case of an incomplete initial database. That is, it may not be reasonable to require that we know that an action is possible in a situation that is incompletely specified. We may prefer to consider explanatory diagnoses, where the theory allows us to consistently assume that the preconditions for $HIST$ or for $HIST \cdot E$ hold, but not necessarily that they are entailed by our theory. To this end, we propose the following alteration to our definition of explanatory diagnoses, Definition 2. A comparable refinement can be made to our definition of assumption-based explanatory diagnosis, Definition 7.

Definition 8 (Potential Explanatory Diagnosis)

A potential explanatory diagnosis for system $(\Sigma, HIST, COMPS, OBS_F)$ is a sequence of actions $E = [\alpha_1, \dots, \alpha_k]$ such that,

- $\Sigma \wedge Poss(HIST \cdot E, S_0)$ is satisfiable, and
- $\Sigma \wedge Poss(HIST \cdot E, S_0) \models OBS_F(do(HIST \cdot E, S_0))$.

Note that $HIST$ is a sequence of actions that we know to have been performed. Thus, the preconditions for each of the actions in $HIST$ are true in the corresponding situations. This provides us with further information concerning the truth values of fluents at various situations, helping to constrain our search.

6.3 VERIFYING LIKELY DIAGNOSES

To further address the problem of generating explanatory diagnoses, we propose exploiting domain information and maintaining a library of (assumption-based) explanatory diagnoses, indexed by observations and/or situation histories. With candidate diagnoses in hand, the problem of computing explanatory diagnoses reduces to a verification problem, rather than a generation problem. Given a system $(\Sigma, HIST, COMPS, OBS_F)$, and a candidate diagnosis E , such that $S = do(HIST \cdot E, S_0)$, we are interested in verifying that E is indeed a diagnosis of the system. Verifying a candidate diagnosis is simply a query evaluation problem. It can be accomplished by regression and theorem proving in the initial database, as per Theorem 1 above.

Example 3

Given the system $(\Sigma, [], \{Pwr, Pmp\}, \neg filling(s))$, and the candidate diagnosis $E = [pwr_failure]$, E can be verified to be an explanatory diagnosis with respect to the system by evaluating the query

$$\begin{aligned} &\mathcal{R}^*[\mathcal{R}_{Poss}[Poss(do(pwr_failure, S_0))]] \\ &\quad \wedge \mathcal{R}^*[\neg filling(do(pwr_failure, S_0))] \end{aligned}$$

with respect to the initial database, Σ_{init} .

7 RELATED WORK

This work has been influenced by formal characterizations of diagnosis for systems without an explicit representation of actions (e.g., (de Kleer et al. 1992), (Reiter 1987), (Console and Torasso 1991)) and by Reiter's work on the frame problem and the problem of temporal projection (Reiter 1992) and (Reiter 1998). Aside from previous work by the author (e.g., (McIlraith 1997), (McIlraith 1997a)), research to date has not explicitly addressed the problem of integrating a rich representation of actions into diagnostic reasoning. As such, there is little related work that exploits a comprehensive representation of action.

A recent notable exception is Thielscher's work on dynamic diagnosis (Thielscher 1997). Thielscher's basic representation is similar in spirit to ours, though he does not use the situation calculus and he does not

exploit an axiomatic solution to the frame, ramification and qualification problems. Like us, he adopts the basic MBD ontology, employs an action theory to represent actions, and represents certain state constraints causally to capture the indirect effects of actions as dictated by the system description of the device. Where he differs is in the actual task he is performing, and the assumptions he makes. To use terminology discussed here, he is computing *what is wrong* diagnoses. Given an action history and some observed aberrant behaviour, Thielscher conjectures that certain components must be abnormal to account for the observed behavior. He is not conjecturing actions. To simplify computation, he assumes that all components are initially normal, and he uses a priori likelihood of failure to select most likely candidate diagnoses. While his paper examines *what is wrong* diagnoses, rather than *what happened* diagnoses, clearly the two are intimately related. (McIlraith 1997) discusses the inter-relationship between these two types of diagnosis in the context of the situation calculus framework described here.

The research on temporal diagnosis and diagnosis of dynamic systems originating in the diagnosis research community (e.g., (Brusconi et al. 1995), (Console et al. 1994), (Hamscher 1991), (Friedrich and Lackinger 1991), (Lackinger and Nejd1 1991), (Dressler 1994)) and in particular (Cordier and Thiébaux 1994) is also loosely related. (Brusconi et al. 1995) recently provided a characterization of temporal abductive diagnosis together with algorithms for computing these diagnoses under certain restrictions. Building on earlier work (Console et al. 1994), they decouple atemporal and temporal diagnoses, using *SD* to represent the behaviour of the atemporal components and transition graphs to represent the temporal components. The later work uses temporal constraints to represent the temporal components. Also related is the work on event-based diagnosis by (Cordier and Thiébaux 1994). Their work is similar in motivation to our work on explanatory diagnosis, viewing the diagnosis task as the determination of the event-history of a system between successive observations. While this work is related, the representation of action is impoverished. It does not provide a comprehensive representation of the preconditions for and the effects of actions, nor does it address the frame, ramification and qualification problems. In particular, their notion of actions is simply an explicit transition systems. This is sufficiently expressive for their application but the lack of a compact representation proves problematic.

In the area of reasoning about action, research on temporal explanation and postdiction has an interesting

relationship to this work (e.g., (Crawford and Etherington 1992), (Baker 1991)). Of particular note is Shanahan's research (Shanahan 1993). While Shanahan also proposes the situation calculus as a representation language for axiomatizing his domain, he does so without an axiomatic solution to the frame and ramification problems. As such these problems must be addressed coincidentally with generating explanatory diagnoses. In contrast, our characterization of explanatory diagnosis, with its axiomatic solution to the frame and ramification problems, enables simpler characterization and computation of temporal explanation.

8 SUMMARY

The results in this paper provide contributions to model-based diagnosis and knowledge representation. Our concern in this paper was, given a system that affects and can be affected by the actions of agents, and given some observed (aberrant) behaviour, how do we capture the notion of *what happened*, i.e., how do we go about conjecturing a sequence of actions that account for the behaviour we have observed.

We addressed this problem by providing a mathematical characterization of the notion of explanatory diagnosis in the context of a rich situation calculus representation, proposed in (McIlraith 1997a). Our characterization made apparent the direct relationship of explanatory diagnosis to the planning task, and in particular to Green's notion of deductive plan synthesis. However, generating explanatory diagnoses is actually akin to planning in the face of state constraints and a potentially incomplete initial database. Our representation scheme addressed the challenges presented by state constraints by providing an axiomatic a priori solution to the frame, ramification and qualification problems. This enabled us to extend Reiter's results on the soundness and completeness of regression and to show that we can generate explanatory diagnoses by regression followed by theorem proving in the initial database. A remaining difficulty was that our initial database is often incomplete. Exploiting features of diagnosis problems, we proposed the notions of assumption-based and potential explanatory diagnosis, to allow for the conjectured sequences of actions that constitute a diagnosis to be predicated on some other assumptions we choose to make about the world. Finally, we proposed exploiting a library of precomputed likely diagnoses, indexed by context and observations. This enabled us to verify, rather than generate explanatory diagnoses.

In our dissertation work (McIlraith 1997), we have also

addressed the complementary problem of conjecturing *what is wrong* diagnoses. In future work we examine the application of GOLOG procedures (Levesque et al. 1997) to diagnostic problem solving.

Acknowledgements

I would like to thank Ray Reiter, Yves Lespérance, Hector Levesque and Allan Jepson for helpful comments on the work presented in this paper. I would also like to thank Oskar Dressler, for engaging me in an interesting discussion regarding this work that helped me look at the material in a different light. Finally, thanks to the reviewers for their thoughtful reviews.

References

- A. Baker (1991). Nonmonotonic Reasoning in the Framework of the Situation Calculus. *Artificial Intelligence* **49**(5-23):361-392.
- V. Brusconi, L. Console, P. Terenziani and D. Theseider Dupré (1995). An Efficient Algorithm for Computing Temporal Abductive Diagnoses. In *Proceedings of the Sixth International Workshop on Principles of Diagnosis*, 41-48.
- L. Console, L. Portinale, D. Theseider Dupré and P. Torasso (1994). Diagnosing Time-Varying Misbehavior: An Approach Based on Model Decomposition. *Annals of Mathematics and Artificial Intelligence* **11**(1-4):381-398.
- L. Console and P. Torasso (1991). A Spectrum of Logical Definitions of Model-Based Diagnosis. *Computational Intelligence* **7**(3):133-141.
- M. Cordier and S. Thiébaux (1994). Event-Based Diagnosis for Evolutive Systems. In *Proceedings of the Fifth International Workshop on Principles of Diagnosis*, 64-69.
- J. Crawford and D. Etherington (1992). Formalizing Reasoning About Change: A Qualitative Reasoning Approach. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, 577-582.
- J. de Kleer (1991). Focusing on Probable Diagnoses. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, 842-848.
- J. de Kleer, A.K. Mackworth and R. Reiter (1992). Characterizing Diagnoses and Systems. *Artificial Intelligence* **56**(2-3):197-222.
- O. Dressler (1994). Model-based Diagnosis on Board: Magellan-MT Inside. In *Proceedings of the Fifth International Workshop on Principles of Diagnosis*, 87-92.
- R. Fikes and N. Nilsson (1971). STRIPS: A New Approach to Theorem Proving in Problem Solving. *Artificial Intelligence* **2**:189-208.
- G. Friedrich and F. Lackinger (1991). Diagnosing Temporal Misbehaviour. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, 1116-1122.
- C. Green (1969). Theorem Proving by Resolution as a Basis for Question-Answering Systems. In B. Meltzer and D. Michie (ed.), *Machine Intelligence 4*, 183-205. New York: American Elsevier.
- W. Hamscher (1991). Modeling Digital Circuits for Troubleshooting. *Artificial Intelligence* **51**(1-3):223-271.
- B. Kramer et al. (1996). Developing an Expert System Technology for Industrial Process Control: An Experience Report. In *Proceedings of the Conference of the Canadian Society for Computational Studies of Intelligence (CSCSI'96)*, 172-186.
- F. Lackinger and W. Nejdl (1991). Integrating Model-Based Monitoring and Diagnosis of Complex Dynamic Systems. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, 1123-1128.
- H. Levesque, R. Reiter, Y. Lespérance, F. Lin and R. Scherl (1997). GOLOG: A Logic Programming Language for Dynamic Domains. In *Journal of Logic Programming, Special Issue on Actions* **31**(1-3):59-83.
- F. Lin and R. Reiter (1994). State Constraints Revisited. *Journal of Logic and Computation* **4**(5):655-678.
- F. Lin and R. Reiter (1995). How to Progress a Database II: The STRIPS Connection. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, 2001-2007.
- J. Lloyd (1987). *Foundations of Logic Programming* (2nd ed.). Springer Verlag.
- S. McIlraith (1994). Further Contributions to Characterizing Diagnosis. *Annals of Mathematics and Artificial Intelligence* **11**(1-4):137-167.
- S. McIlraith (1994a). Towards a Theory of Diagnosis, Testing and Repair. In *Proceedings of The Fifth International Workshop on Principles of Diagnosis*, 47-53.
- S. McIlraith (1997). Towards a Formal Account of Diagnostic Problem Solving. PhD thesis, Department of Computer Science, University of Toronto, Toronto, Ontario, Canada.

S. McIlraith (1997a). Representing Actions and State Constraints in Model-Based Diagnosis. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97)*, 43–49.

D. Poole (1988). Representing Knowledge for Logic-Based Diagnosis. In *Proceedings of the Fifth Generation Computer System Conference (FGCS-88)*, 1282–1290.

R. Reiter (1987). A Theory of Diagnosis from First Principles. *Artificial Intelligence* **32**:57–95.

R. Reiter (1991). The Frame Problem in the Situation Calculus: A Simple Solution (sometimes) and a Completeness Result for Goal Regression. In *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of J. McCarthy*, 359–380. San Diego, CA: Academic Press.

R. Reiter (1992). The Projection Problem in the Situation Calculus: a Soundness and Completeness Result, with an Application to Database Updates. In *Proceedings First International Conference on AI Planning Systems*, 198–203.

R. Reiter (1996). Natural Actions, Concurrency and Continuous Time in the Situation Calculus. In L. Aiello, J. Doyle and S. Shapiro (ed.) *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR'96)*, 2–13. Cambridge, Massachusetts: Morgan Kaufmann.

R. Reiter (1998). *KNOWLEDGE IN ACTION: Logical Foundations for Describing and Implementing Dynamical Systems*. In preparation.

M. Shanahan (1993). Explanation in the Situation Calculus. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, 160–165.

M. Shanahan (1997). Event Calculus Planning Revisited. In *Proceedings 1997 European Conference on Planning (ECP 97)*, Springer-Verlag Lecture Notes in Artificial Intelligence (no. 1348), 390–402.

M. Thielscher (1997). A Theory of Dynamic Diagnosis. In *Linköping Electronic Articles in Computer and Information Science* **2**(11). Research article received for discussion October, 1997. Revised March, 1998. <http://www.ep.liu.se/ea/cis/1997/011/>.

R. Waldinger (1977). Achieving Several Goals Simultaneously In E. Elcock and D. Michie (ed.) *Machine Intelligence* 8,94–136. Edinburgh, Scotland: Ellis Horwood.