A Formal Account of Nondeterministic and Failed Actions

James P. Delgrande
School of Computing Science
Simon Fraser University
Burnaby, B.C. V5A 1S6 Canada
jim@cs.sfu.ca

Hector J. Levesque
Dept. of Computer Science
University of Toronto
Toronto, ON M5G 3G4 Canada
hector@ai.toronto.edu

Abstract

Nondeterminism is pervasive in all but the simplest action domains: an agent may flip a coin or pick up a different object than intended, or an action may fail and may fail in different ways. In this paper we provide a qualitative theory of nondeterminism. The account is based on an epistemic extension to the situation calculus that accommodates sensing actions. Our position is that nondeterminism is an epistemic phenomenon, and that the world is most usefully regarded as deterministic. Nondeterminism arises from an agent's limited awareness and perception. The account offers several advantages: an agent has a set of categorical (as opposed to probabilistic) beliefs, yet can deal with equally-likely outcomes (such as in flipping a fair coin) or with outcomes of differing plausibility (such as an action that may on rare occasion fail).

1 Introduction

In standard accounts of reasoning about action, an agent (or agents) carries out a sequence of actions, possibly including sensing, where each action has a fixed, determined (though possibly conditional) effect on the environment. The real world of course is nothing like this. Not only will an agent have incomplete information, it may also have inaccurate or incorrect information. Actions may fail, perhaps because a prerequisite condition is not satisfied, or an action may simply fail with no known reason. Moreover, the outcomes of some actions, such as flipping a coin, may be impossible to predict. Such actions in which the outcome is not fully predictable are called *nondeterministic*, in contrast to the traditional account, which concerns *deterministic* actions. Nondeterminism in reasoning about action has often been addressed via a *quantitative* account based on probability theory. In such an account, action outcomes are given by a (known) probability distribution over the different ways that the world may evolve.

Our goal in this paper is to provide a *qualitative* account of nondeterminism. That is, we want to be able to say that flipping a coin will have it come out heads or tails (and not heads with probability 0.5 and the same for tails). We also want to be able to state that pressing a light switch may fail, but that the agent will disregard this possibility unless further evidence (e.g., sensing that the room remains dark) indicates otherwise. There are several advantages to a qualitative account of nondeterminism. First, it is *intuitive*; people in general assume that a light switched on will go on, and don't consider the alternatives unless there is a need to. A qualitative approach is *simpler* than a probabilistic account, since one does not need to manipulate probabilities. A qualitative account also allows an agent to maintain a knowledge base of categorical beliefs, rather than assertions with probabilities. Consequently it may also offer computational advantages. A common criticism of probabilistic accounts is that it is not clear where the numbers come from, nor is it always easy to judge whether a given probability is reasonable or not. On the other hand, probabilistic accounts have proven to be highly useful in many applications. Hence, any qualitative approach should be *compatible* with a probabilistic approach (which we argue is the case here). Overall then, we see a qualitative approach as having independent interest and utility, and complementing numeric approaches.

So what is qualitative nondeterminism? Our view is that it is an *epistemic* notion, reflecting an agent's limited awareness of the domain and how it might evolve. That is, as in the traditional account of reasoning about action, we assume that the world is deterministic and that each state of the world is determined by its predecessor. Thus, in flipping a coin, an agent in fact executes an action that will correspond to a flip-heads or a flip-tails action; however it will not know which, at least until after the coin is observed. Since our account allows action outcomes with different plausibility levels, the notion of a *failed action* is subsumed by our approach, in that a failed action is an unlikely or implausible nondeterministic outcome of an action execution. This view of the world as being deterministic is usually associated with Laplace [Laplace, 1814], as the initial proponent. While this view of the universe is incorrect, at least if one accepts quantum theory, at the level of commonsense reasoning it provides a highly useful approximation in the context of reasoning about action.

Given our assumption of a deterministic world, a sufficiently knowledgeable and perceptive agent would know whether an unseen coin toss yields heads or tails. Our agent, though, isn't sufficiently perceptive, so all it knows is there will be one of the two possible outcomes. We formalise this notion by making use of an epistemic extension to the situation calculus, and broadly building on [Bacchus *et al.*, 1999] and [Delgrande and Levesque, 2012]. In common with these papers, we express that an agent intending to execute an action may inadvertently execute another. We also attach plausibility values to such alternative actions. Hence in pressing a light switch, the agent may be very likely to press the switch but may, implausibly, press the wrong switch. In the approach, the agent will believe that it has pressed the switch unless further evidence compels it to believe otherwise.

Arguably, the approach provides an intuitive, declarative, qualitative approach to nondeterminism, including actions with equally-plausible outcomes, as well as actions with implausible outcomes. It allows for the specification of domain-specific theories of types of action failure, thereby also shedding light on the formal underpinnings of reasoning about action. The next section presents background material, while the following section discusses the approach. After this we present the formal details of the approach, and explore its properties and application. This is followed by a discussion and a conclusion.

2 Background

2.1 The Situation Calculus

We describe the version of the situation calculus presented in [Levesque *et al.*, 1998; Reiter, 2001]. The language of the situation calculus is first-order with equality, with sorts for actions, situations, and objects (everything else). A *situation* represents a world history described by a sequence of actions. Predicates and functions whose values are situation-dependent (and whose last argument is a situation) are called *fluents*. A set of initial situations expresses the ways the domain might be initially, while the constant S_0 denotes the actual initial state of the domain. The term $do(a, s)$ denotes the unique situation that results from executing action a in situation s. The notation $\langle a_1, \ldots, a_n \rangle$ specifies an action sequence. This is also used to abbreviate the term $do(a_n, do(\ldots, do(a_1, s) \ldots))$ as: $do(\langle a_1, \ldots, a_n \rangle, s)$.

To axiomatise a dynamic domain, we use *basic action theories* [Reiter, 2001] consisting of (1) axioms which describe the initial states of the domain, including the initial beliefs of the agents; (2) precondition axioms, giving the conditions under which each action can be executed;[1] (3) successor state axioms, which describe how each fluent changes as the result of actions; (4) sensing axioms for each action, described below; (5) unique names axioms for the actions; and (6) domain-independent foundational axioms.

[Scherl and Levesque, 2003] axiomatises an agent's knowledge by treating situations as possible worlds. Two special fluents were used, *SF* and B. An action returns a (binary) sensing result, and $SF(a, s)$ holds when action a returns sensing value 1 in situation s. Sensing axioms give the conditions under which $SF(a, s)$ holds. The B fluent is the usual belief accessibility relation: $B(s', s)$ holds when the agent in situation s thinks that situation s' might be the actual situation.[2] A successor state axiom for B is given by:[3]

$$B(s', do(a, s)) \equiv \quad (1)$$
$$\exists s^*[B(s^*, s) \wedge s' = do(a, s^*) \wedge (SF(a, s^*) \equiv SF(a, s))].$$

So the situations s' that are B-related to $do(a, s)$ are the ones that result from doing action a in a previously related situation s^*, such that the sensor associated with action a has the same value in s^* as it does in s. Belief then is defined, as usual, as truth in all accessible situations:[4]

$$Bel(\phi, s) \doteq \forall s'.B(s', s) \supset \phi[s'].$$

2.2 Nondeterminism in Reasoning About Action

Most work in nondeterminism assumes that an action may have indeterminate consequences: in a quantitative account, the effect of an action is given by a probability distribution over possible next states; in a qualitative account, the effect of an action is a set of possible next states.

For quantitative approaches, in the planning community, stochastic domains are often modelled using Markov Decision Processes [Puterman, 1994]; see for example [Cassandra *et al.*, 1994]. In a different vein, Poole's Independent Choice Logic [Poole, 2008] adds independent stochastic inputs to an acyclic logic program. A probability distribution is given over each set of alternatives, and states of the world are assigned a probability derived from these distributions.

[Bacchus *et al.*, 1999] introduces epistemic aspects to nondeterminism. A fluent *OI*, for *observational indistinguishability*, is introduced where $OI(a, a's)$ expresses that execution of actions a and a' in situation s cannot be distinguished by the agent. Actions a and a' may be "noisy"; for example, $advance(3.0, 2.8)$ might be the action where an agent attempting to move 3.0 units actually moves 2.8 units. Probabilities are assigned to the initial situations, and these probabilities are adjusted appropriately following the execution of actions.

In qualitative approaches, nondeterminism can be modelled by *transition systems*, that is, labelled, directed graphs where vertices represent states of the world, and arcs are labelled by actions. Nondeterminism is expressed by having more than one arc leaving a vertex with the same label. Transition systems provide a semantic basis for *action languages* [Gelfond and Lifschitz, 1998]; see e.g. [Baral, 1995] for an account of (among other issues) nondeterminism. Similar work has been carried out in the planning community: for example, in [Cimatti *et al.*, 2003] an action is associated with a set of possible next states; [Jensen *et al.*, 2004] treats actions as having an expected outcome along with a set of possible unexpected outcomes (for example, failing).

In the circumscriptive event calculus [Shanahan, 1997], actions with nondeterministic effects depend on *determining fluents* whose value is unknown. Hence this can be seen as

[1] For simplicity, we will assume all actions are always possible.

[2] Because B is a fluent, the order of the situation arguments is reversed from the usual convention in modal logic.

[3] Free variables are assumed to be universally quantified from outside.

[4] The formula ϕ here is assumed to have a single free variable of sort situation. It is usually written either with a distinguished variable now as a placeholder, or else the variable is suppressed. Either way, $\phi[t]$ denotes the formula with the variable replaced by t.

"nature" influencing the outcome of an action. Another approach to nondeterminism is provided by the Golog family of languages [Levesque *et al.*, 1997]). Nondeterminism there is handled as a programming construct, like in dynamic logic [Pratt, 1976]: every action is considered to be a program, and given two programs α and β, the program $[\alpha|\beta]$ is the nondeterministic choice of one of them. This idea is refined in [Boutilier *et al.*, 2000] where $BestDo(\alpha, s, s')$ holds when the final outcome maximizes utility.

[Delgrande and Levesque, 2012] addresses revision and fallible actions in the situation calculus. There, an agent intending to execute one action may inadvertently execute another. This is expressed by a predicate Alt, where $Alt(a, a', s)$ asserts that an agent intending to execute a may, implausibly, execute a' in s. Intended actions are believed to have been executed, while alternative actions are taken as being maximally implausible.

3 The Approach

As stated, our goal is to provide a logical theory of nondeterminism. Moreover, we take a Laplacian stance: the world is in fact deterministic, and nondeterminism is an artifact of the agent's limited senses and limited knowledge. Examples of nondeterminism that we account for include the following:

1. An agent flips a fair coin. Afterwards it believes that the coin landed heads or tails. Perhaps also, the agent believes that it is just possible that the flipping action failed, and it is extremely implausible that the coin landed on its edge.

2. An agent throws a dart at a dartboard. Since it is a poor player, it believes that it is equally plausible that the dart is embedded in the dartboard or in the surrounding cork.

3. An agent intends to push a button. Afterwards it believes that it pushed the button, but believes it is possible (but implausible) that it pushed a nearby button.

4. An agent knows that a light is on, and toggles the light switch twice. On observing that the light is off, it concludes that one of the toggle actions failed.

Consider the first example, flipping a coin, but with outcome heads or tails only. Call these actions f, for $flip$, and fH and fT, for flipping heads and tails respectively. To begin, there are two things that we do *not* mean by nondeterminism. First, there is no notion of *causality*; we do not want to say that a f action causes a fH. As well, we don't want to say that f is a composite action, made up of a nondeterministic disjunction of fH and fT. As an *account* of nondeterminism, such a specification is of course circular.

Arguably, f is a "real" action, to the extent that one can talk about an agent flipping a coin, and it makes sense to talk about what happens when a coin is flipped. Now, when an agent attempts to execute a f action, either a fH or fT is executed; an agent would know this but wouldn't know *which* is executed. We call an action such as f a *virtual action*, by analogy with virtual memory. I.e., from a user's point of view, virtual memory behaves just like physical memory in that it has addresses at which data is stored; however, where the data is *really* stored is at some physical location (with its own address) and beyond the control of the user.

Given that we take nondeterminism to be an epistemic phenomenon, we focus on the evolution of an agent's beliefs. For simplicity we assume a single agent. An agent may carry out a physical action or it may sense its environment. Physical actions may have nondeterministic consequences. Due to space limitations, we assume that sensing is accurate, in that an agent will believe the result returned by a sensor, though maintaining the possibility that the sensor in fact failed in some way.

Fundamental to our account is the notion of *plausibility*, where a plausibility value is a non-negative integer. Plausibilities are used in two related ways. First, an agent's belief state will be modelled by a plausibility ordering ([Spohn, 1988; Darwiche and Pearl, 1997]) over a set of situations. That is, situations are assigned plausibility values, where a lower value means that a situation is considered to be more plausible. Situations with plausibility 0 characterize the agent's categorical beliefs. Second, ordered pairs of actions may have an associated plausibility in a situation. If plausibility p is associated with actions a_i and a, this means that if the agent intends to execute a_i, then with plausibility p action a may be executed. For example, for a f action, fH and fT will each have plausibility 0, and a failed action, $null$ may have plausibility 1. These values are interpreted to be a qualitative measure, where non-zero values capture "big step probabilities" or large, orders-of-magnitude values.

Our goal then is to track the evolution of an agent's belief state, expressed as a plausibility ordering, following physical and sensing actions. However, nondeterminism complicates matters. Consider again flipping a coin. Following a f, the world will evolve so that the coin lands heads, or so that it lands tails. However the agent will believe only that the coin shows heads or it shows tails. Consequently, when there are nondeterministic actions, the actual physical actions that occur are not enough to determine the situations that the agent considers possible; we need to also consider what the agent *believes* what the outcomes may be. This means that in tracking the evolution of an agent's belief state, we will need to keep track of both those actions that actually were executed, along with those actions that the agent intended to execute. See [Delgrande and Levesque, 2012] for more on this.

4 The Formal Framework

In this section we describe the new theory of reasoning about and with nondeterministic actions. We begin with a *basic action theory*, as summarised in Section 2. Due to space limitations we omit the BAT axioms; however, see [Levesque *et al.*, 1998]. To a basic action theory, we add two predicates, Alt, for describing alternative actions and B, for capturing an agent's belief state. These predicates generalise similarly-named predicates in [Delgrande and Levesque, 2012], which in turn generalises [Bacchus *et al.*, 1999; Shapiro *et al.*, 2011]; see Section 6 for a discussion.

We express the fact that an agent intending to execute (physical) action a_1 may in fact execute a_2 with plausibility p in situation s by the formula

$$Alt(a_1, a_2, p, s).$$

The plausibility p is a qualitative measure. If $p = 0$ then a_1

and a_2 are equally plausible. If $p = 1$ then a_2 is implausible wrt a_1, and a_2 would be a surprise, should the agent learn that a_2 was in fact executed. Higher values of p are also possible. This then captures both qualitative nondeterminism (for $p = 0$) as well as unintended or unlikely actions (for $p > 0$).

Alt is not a "regular" fluent in that it does not have a successor state axiom; rather Alt is defined in terms of other fluents. (Thus Alt behaves much like $Poss$ in basic action theories.) For each action we assume an axiom of the following form, where the free variables of A are among $\vec{x}$, and Ψ_A is uniform in s[5] [Levesque *et al.*, 1998] and does not mention Alt or B:

Alt Axiom: $Alt(A(\vec{x}), a, p, s) \equiv \Psi_A(\vec{x}, a, p, s)$

Consider how our earlier examples may be expressed:

1. Flipping a coin: Let f be the flipping action and fH and fT flipping heads and tails respectively:

   $Alt(f, a, p, s) \equiv (a = fH \wedge p = 0) \vee (a = fT \wedge p = 0)$

   f is a virtual action, and attempting to execute it results in fH or fT. To say that the flipping action might fail, or that the coin might land on its edge, one could disjoin:

   $$(a = null \wedge p = 1) \vee (a = fEdge \wedge p = 2)$$

2. Throwing a dart: Let tB be the action of throwing a dart so it hits the dartboard; and tW the action where the dart hits the adjacent wall. If the light isn't dim, the dart usually hits the board; if it is, it may also hit the wall.

   $$\begin{aligned} &Alt(tB, a, p, s) \equiv \\ &\quad \neg Dim(s) \supset (a{=}tB \wedge p{=}0) \vee (a{=}tW \wedge p{=}1)) \quad \wedge \\ &\quad\quad Dim(s) \supset ((a{=}tB \wedge p{=}0) \vee (a{=}tW \wedge p{=}0)) \end{aligned}$$

3. Pushing a light switch:

   $Alt(push(x), a, p, s) \equiv (a = push(y) \wedge p = |\, x - y \,|)$

4. Toggling (t) a light switch where the toggling may fail, implausibly and for no known reason:

   $Alt(a, a', p, s) \equiv (a = a' \wedge p = 0) \vee (a = null \wedge p = 1)$

Consider next the evolution of an agent's belief state. A belief state is made up of a set of situations along with a plausibility value attached to each situation. Situations with plausibility 0 are those that the agent considers may correspond to the real world; situations with plausibility > 0 provide (relatively) plausible alternatives, should the agent learn that its beliefs are incorrect. We need to track such a belief state given (1) a finite sequence of actions σ that at each point in time the agent intended to execute, and (2) a situation representing the actual state of affairs. This is expressed using a 4-place fluent

$$B(s', p, \sigma, s)$$

where for specific values of σ and s, $B(\cdot, \cdot, \sigma, s)$ characterises a belief state.

It is straightforward to describe an initial belief state. An initial situation is defined to be one with no predecessor:

$$Init(s) \doteq \neg\exists a, s'.\, s = do(a, s')$$

[5] Roughly, the fluents in Ψ_A are with respect to s only.

The following axioms describe initial situations; S_0 by convention represents the initial state of the actual situation. Recall that an action sequence is represented as a list of actions surrounded by $\langle \ldots \rangle$. We assume some appropriate axiomatisation for lists; the concatenation of a to σ is given by $\sigma \cdot a$.

Axioms:

1. $Init(S_0)$
2. $Init(s) \supset \exists n B(s, n, \langle\,\rangle, S_0)$
3. $B(s', n, \sigma, S_0) \supset Init(s') \wedge \sigma = \langle\,\rangle$
4. $B(s', n_1, \langle\,\rangle, S_0) \wedge B(s', n_2, \langle\,\rangle, S_0) \supset n_1 = n_2$

Axiom 2 asserts that every initial situation has some plausibility according to the agent. Axiom 3 asserts that only initial situations are considered possible by the agent at S_0, and the agent believes that it has attempted no actions. The fourth axiom states that at the initial "index" $\langle\,\rangle$,S_0, plausibility is a function of a given situation.

An agent may change its beliefs by ***sensing*** its environment or by (attempting to) *execute* a physical action. For convenience, and unlike [Scherl and Levesque, 2003], sensing and physical actions are two distinct action types; thus we have the axioms, with obvious intended interpretation:

Axiom 5: a. $Action(a) \equiv (PhysAct(a) \vee SenseAct(a))$
b. $\neg(PhysAct(a) \wedge SenseAct(a))$

Consider physical actions, and where the agent intends to execute action a_i. Two things need to be tracked. First the agent may intend a_i but in fact executes a (thus a light switch other than that intended may be pushed). As discussed, this is kept track of in the last two arguments of the B fluent. But second, the agent will also be aware of potential alternative actions. Hence, in pushing a light switch, the agent will believe that it has pressed the correct switch, but it also knows that it may possibly press the incorrect switch. Hence such alternatives also need to be kept track of. Specifically, consider where $B(s', p, \sigma, s)$ is true, and so in the belief state indexed by σ, s, situation s' has plausibility p. If the agent intends to execute a_i where $Alt(a_i, a, p', s)$ is true, then the situation $do(a, s')$ should be assigned plausibility $p + p'$. Thus, if $B(s', p, \sigma, s)$ and $Alt(a_i, a_i, 0, s)$ are true, then for intended action a_i, the plausibility of situation $do(a_i, s')$ will be unchanged. If $Alt(a_i, a_j, 1, s)$ is true, then the plausibility of situation $do(a_j, s')$ will be $p + 1$; i.e. under an unexpected action the resulting situation will become less plausible.

For sensing, we adopt a simple model wherein the agent believes the result of sensing, but also allows that (implausibly) sensing may be incorrect. Plausibilities are modified following the recipe of [Darwiche and Pearl, 1997]: If g is sensed to be true, then the plausibility of all accessible B-related situations where g is false is increased by 1. Thus all $\neg g$-situations are believed to be less plausible. Also, the plausibility of situations in which g is true is uniformly reduced so that some situation in which g is true has plausibility 0. As a result, the agent believes g after sensing that it is true.

Two definitions are introduced to assist in the expression of the successor state axiom for B. The first captures that the sensing action a has the same result in situations s and s':

$$Agree(a, s', s) \;\doteq\; \mathit{SF}(a, s') \equiv \mathit{SF}(a, s).$$

Recall from Section 2 that a is a sensing action associated with some fluent; *SF* is the result of sensing that fluent in s.

The second abbreviation asserts that the minimum plausibility of ϕ at some belief state is d.

$$Min(\phi, d, \sigma, s) \doteq \exists s' [B(s', d, \sigma, s) \wedge \phi[s'] \wedge \\ \forall s'', d'((d' < d \wedge B(s'', d', \sigma, s)) \supset \neg\phi[s''])]$$

The successor-state axiom for B can be given as follows:

Axiom 6:

$$\begin{aligned} & B(s', n, \sigma, do(a, s)) \equiv \\ & \quad \exists s^*, n^*, a^*, \sigma^*, a_i, p_2. \\ & \quad\quad B(s^*, n^*, \sigma^*, s) \wedge s' = do(a^*, s^*) \wedge \sigma = \sigma^* \cdot a_i \wedge \\ & \quad\quad\quad \exists p_1 Alt(a_i, a, p_1, s) \wedge Alt(a_i, a^*, p_2, s^*) \wedge & (2) \\ & \quad\quad [PhysAct(a) \wedge n = n^* + p_2] \vee & (3) \\ & \quad\quad [SenseAct(a) \wedge a = a_i \wedge a = a^* \\ & \quad\quad\quad ((Agree(a,s,s^*) \wedge n = n^* - d \wedge \\ & \quad\quad\quad\quad Min(Agree(a,s,s^*),d,\sigma,s)) & (4) \\ & \quad\quad \vee (\neg Agree(a,s,s^*) \wedge n = n^* + 1))] & (5) \end{aligned}$$

The axiom gives the conditions for situation s' to have plausibility n in the belief state indexed by σ and $do(a, s)$. The right-hand side of the equivalence is made up of three parts, a preamble (2), and parts for physical (3) and sensing (4,5) actions. For the preamble (2), there is an intended action a_i and action sequence σ^* where $\sigma = \sigma^* \cdot a_i$ and $Alt(a_i, a, p_1, s)$ holds. That is, a_i is intended but a is actually executed. As well, s' is the result of an action a^* where $Alt(a_i, a^*, p_2, s^*)$ holds. So the agent intends to execute action a_i but in fact executes a. Action a^* is also an alternative to a_i.

For a physical action (3), the agent has no access to the action a that was actually executed. Action a_i was intended, and so for alternative action a^*, the resulting (i.e. $s' = do(a^*, s^*)$) situation would have plausibility increased by p_2. For the "normal" case we would have $a^* = a_i$ and $p_2 = 0$.

Sensing connects the agent with the real situation. The effect of a sensing action is to, if necessary, uniformly decrease the plausibility value of those situations in which the sensing result agrees with that at s, such that some resulting situation has plausibility 0 (4), and to increase the plausibility of other situations by 1 (5).

This characterises how an agent's belief state evolves given an intended and actually-executed action. The agent's categorical beliefs are characterised by 0-ranked situations:

$$Bel(\phi, \sigma, s) \doteq \forall s'. B(s', 0, \sigma, s) \supset \phi[s']. \quad (6)$$

5 Properties of the Framework

Σ below will denote a basic action theory (BAT) as given in the previous section. We use $s' \sqsubseteq s$ to express that s' is a subhistory of s. First, B can be interpreted as a plausibility ordering indexed by the last two arguments of B:

Theorem 1 $\Sigma \models (B(s', p_1, \sigma, s) \wedge B(s', p_2, \sigma, s)) \supset p_1 = p_2$

Ignoring nondeterminism, action progression reduces to that of [Scherl and Levesque, 2003] (SL) for knowledge:

Theorem 2 *Let Σ be a BAT entailing $B(S_0, 0, \langle\rangle, S_0)$ and with sole Alt axiom: $Alt(a, a', p, s) \equiv a = a' \wedge p = 0$. Then there is a SL BAT*[6] *Σ' s.t. for any σ and $s = do(\sigma, S_0)$:*

$$\Sigma \models B(s', 0, \sigma, s) \textit{ iff } \Sigma' \models B(s', s)$$

[6] I.e. Using (1) for B and not mentioning Alt.

Assume every fluent has a corresponding sensing action, and that for any fluent formula ϕ there is a corresponding sensing action $sense_\phi$ defined in the obvious fashion. Then, a revision operator can be defined in terms of sensing:

Theorem 3 *Define, in the case where $\Sigma \models \phi(S_0)$:*

$$BS(\Sigma * \phi) \doteq \{\psi | \Sigma \models Bel(\psi, \langle sense_\phi \rangle, do(sense_\phi, S_0))\}$$

Then (where defined) $BS(\Sigma * \phi)$ satisfies the AGM revision postulates [Gärdenfors, 1988].

A *virtual action* was defined as an action a_v such that no instance of $Alt(a_i, a_v, p, s)$ is entailed. The last result shows that such actions play no role in the evolution of a situation indexing an agent's belief state, nor in the situations characterizing the belief state:

Theorem 4 *Let a_v be a virtual action wrt Σ and let $\Sigma \models B(s', p, \sigma, s)$. Then s', s do not mention a_v.*

We next go through two extended examples, to show the framework in action.

Flipping a Coin There is just one fluent H indicating that the coin is showing heads; tails is given by the negation of H.

The four actions, f, fH, fT, $null$, represent a flip, flip heads, flip tails, and the null action respectively. There is a sensing action, sH for *sense heads*. S_0 is the sole initial situation. The basic action theory is given by the following:

$$\begin{aligned} & H(S_0), \qquad B(S_0, 0, \langle\rangle, S_0) \\ & H(do(a, s)) \equiv (a = fH \vee (a = null \wedge H(s))) \\ & Alt(f, a, p, s) \equiv (a = fH \wedge p = 0) \vee \\ & \quad (a = fT \wedge p = 0) \vee (a = null \wedge p = 1) \\ & SF(a, s) \equiv H(s) \vee a \neq sH \end{aligned}$$

The following are entailed in the case where f is intended and fH is actually executed:

$$\begin{aligned} & B(do(fH, S_0), 0, \langle f \rangle, do(fH, S_0)), \\ & B(do(fT, S_0), 0, \langle f \rangle, do(fH, S_0)), \\ & B(do(null, S_0), 1, \langle f \rangle, do(fH, S_0)) \end{aligned}$$

The agent believes it executed a fH or fT action. The agent admits the possibility that the flip action failed, but this possibility isn't part of the agent's set of beliefs. Following a sensing action sH, the result of the flip can be learned. This is all illustrated in Figure 1.

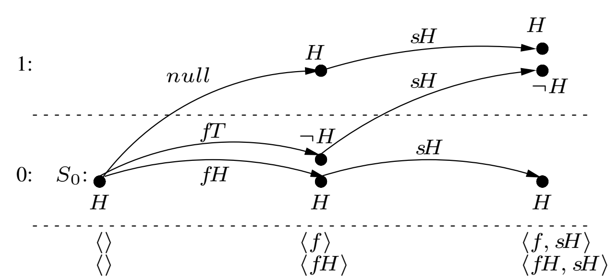

Figure 1

The bottom of the figure shows the evolution of two sequences of actions, the agent's intended actions above, and below them the evolution of the actual situation. Each black dot represents a situation and each such situation lies in a numbered stratum, giving that situation's plausibility. Thus a "column" of black dots, or situations, represents an instance

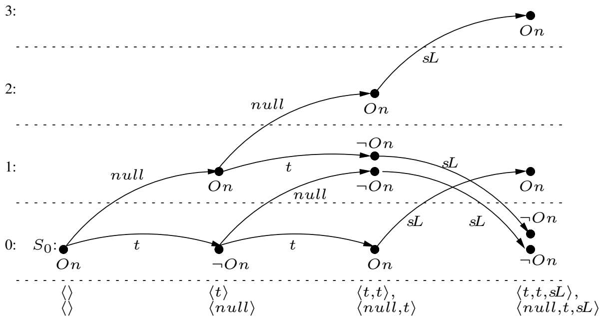

Figure 2

of the agent's belief state. Finally, arcs represent situation transitions, labelled by the appropriate action.

Thus we see that initially (as illustrated by the left-most dot) S_0 is the only possible situation, according to the agent. Following an intended f, but where fH was actually executed, the agent believes that either a fH or fT was executed (plausibility 0) or, implausibly, the action failed (plausibility 1). Last, the sensing action yields the true state of affairs to the agent. Consequently we also have that $\neg Bel(H, \langle f\rangle, do(fH, S_0)$ but $Bel(H, \langle f, sH\rangle, do(sh, do(fH, S_0))$ are entailed.

Toggling a Light Switch A light is initially *on*, and this is known by the agent. Toggling the switch changes the state of the light from *on* to *off* or vice versa. The agent toggles the light switch twice, and observes that the light is *off*. It concludes only that one of the toggling actions must have failed.

There is one fluent On indicating that the light is *on*. There are three actions, t, sL, and $null$ which toggle the light switch, sense the light, and do nothing, respectively. There is one initial state S_0 in which the light is On. The basic action theory is given as follows:

$$\begin{array}{ll} On(S_0), & B(S_0, 0, \langle\rangle, S_0) \\ \multicolumn{2}{l}{On(do(a,s)) \equiv (a = t \wedge \neg On(s)) \vee (a \neq t \wedge On(s))} \\ \multicolumn{2}{l}{SF(a,s) \equiv On(s) \vee a \neq sL} \\ \multicolumn{2}{l}{Alt(t,a,p,s) \equiv (a = t \wedge p = 0) \vee (a = null \wedge p = 1)} \end{array}$$

The agent intends to execute two toggle actions; thus for any $x, y \in \{t, null\}$, $B(do(\langle t,t\rangle, S_0), 0, \langle t,t\rangle, do(\langle x,y\rangle, S_0))$ is entailed.[7] More perspicuously, we obtain

$$Bel(On, \langle t,t\rangle, do(\langle x,y\rangle, S_0))$$

The agent believes that the light is On, regardless of whether toggles were or were not in fact executed.

Further, for any $x, y \in \{t, null\}$ the following are entailed.

$$\begin{array}{l} B(do(\langle null, t\rangle, S_0), 1, \langle t,t\rangle, do(\langle x,y\rangle, S_0)), \\ B(do(\langle t, null\rangle, S_0), 1, \langle t,t\rangle, do(\langle x,y\rangle, S_0)), \\ B(do(\langle null, null\rangle, S_0), 2, \langle t,t\rangle, do(\langle x,y\rangle, S_0)) \end{array}$$

Thus the agent allows for the unlikely possibility that a toggle failed, and even more implausibly, that it failed twice.

Next consider the specific situation where the agent executes two toggle actions, and observes that the light is not On. Assume that in fact the first *toggle* action failed. This is illustrated by Figure 2.

[7] Recall that $do(\langle x,y\rangle, S_0))$ abbreviates $do(y, do(x, S_0))$.

Again, an agent's belief state is indicated by a column of black dots, divided into strata which give the plausibility of each situation, and with the index of that belief state given below each column of dots. It can be seen that for the first two toggle actions, the agent believes that the toggling succeeds, even though the first in fact fails. Also, the agent considers it implausible (with plausibility 1) that one of the toggles failed, and even more implausible (plausibility 2) that both failed.

The agent's incorrect beliefs get sorted out following the sL sensing action. It is discovered that the light is in fact not On; consequently the situations where $\neg On$ is true are "moved down" in the plausibility ranking so that the minimum $\neg On$ situations have plausibility 0. As well, the plausibility of On-situations is increased by 1.

More formally, we have that the following are entailed:

$$B(do(\langle null, t, sL\rangle, S_0), 0, \langle t, t, sL\rangle, do(\langle null, t, sL\rangle, S_0)) \quad (7)$$
$$B(do(\langle t, null, sL\rangle, S_0), 0, \langle t, t, sL\rangle, do(\langle null, t, sL\rangle, S_0)) \quad (8)$$

As a result, the agent knows that one of the toggles – but not which – failed. Hence we do not obtain *recency*, that is, where more recent actions are assumed to be more plausible. This is in contrast with most work in belief revision, where a more recent formula for revision takes precedence over an earlier one, if the formulas conflict (see e.g. [Peppas, 2008]).

Finally, up to this point, the agent's initial belief state was characterized by the situation S_0, and so it is *inconceivable* to the agent that the light is initially not On. Assume there is a second initial situation S_1, and that $\neg On(S_1)$ and $B(S_1, 1, \langle\rangle, S_0)$ hold. Thus the agent believes that the light is On, but that it is not *impossible* (just *implausible*) that the light is not On. If now the agent intends to toggle the light switch twice and the first toggle fails, we obtain in addition to (7), (8) that the following is entailed:

$$B(do(\langle t, t, sL\rangle, S_1), 0, \langle t, t, sL\rangle, do(\langle null, t, sL\rangle, S_0))$$

Thus, the agent allows for an additional possibility where the toggles succeeded, but the light was not initially On. Finally if it were deemed that it was *highly* unlikely that the light was *off*, and so S_1 had plausibility > 1, then the agent would again believe only that a toggle had failed.

6 Discussion

Related Work One branch of work involving nondeterminism makes use of probabilities. The closest such work to the present is [Bacchus *et al.*, 1999]. Somewhat analogous to our use of Alt with plausibility 0, they employ, first, "noisy actions" defined in terms of Golog constructs, and second an "observation indistinguishability" predicate to assert when an agent cannot distinguish one action from another. In addition, quantitative information is added by introducing fluents that assign probabilities to (accessible) situations and to actions. Given that the approach deals with nondeterminism via probabilities, it would not handle the case where an agent turns a light *on*, observes that it is *off*, and concludes categorically that the switching action must have failed.

The closest work to the present is [Delgrande and Levesque, 2012], although the focus there was on belief revision, which we do not consider here. The Alt predicate is used, but in a restricted setting: Intended actions have plausibility 0, and unintended actions are assigned plausibilities higher than any existing plausibility; as a result, *recency* is obtained. In contrast, one may have an arbitrary plausibility relation between intended and actual actions here, and recency is not obtained. Last, our treatment of physical actions in the B fluent is simpler than that of [Delgrande and Levesque, 2012], while also significantly more general.

Other work addressing nondeterminism, including [Baral, 1995; Shanahan, 1997; Levesque *et al.*, 1997; Giunchiglia *et al.*, 2004], generally assumes that the world itself is nondeterministic. As well, nondeterministic outcomes are equally likely and there is no incorporation of plausibility.

The Qualification Problem The qualification problem [McCarthy, 1977] as regards actions refers to the fact that in a realistic setting, it is impossible to list all preconditions for an action. Solutions to the problem have used meta-level considerations (e.g. [Ginsberg and Smith, 1988]), or employed nonmonotonic formalisms (e.g. [Thielscher, 2001]). The present approach arguably resolves (or at least accommodates) the qualification problem, and within first-order logic: Actions are expressed with respect to their "usual" preconditions. As well, using Alt, known ways in which an action may fail or have unusual effects can be encoded. Further, the evolution of an agent's beliefs can allow for the fact that (implausibly) an action may fail and indeed may fail for no known reason. The fact that an action failed or had an unexpected outcome is something that would presumably be discovered by subsequent sensing.

Relation with Probabilistic Approaches One of the goals of the approach was that it be "compatible" with probabilistic approaches. Arguably the approach *is* compatible. Specifically, there is nothing about the plausibility arguments in Alt and B that says that they have to refer to plausibilities. Thus, $Alt(a_1, a_2, p, s)$ could have the intended reading that an agent meaning to execute a_1 will execute a_2 with probability p in situation s. Then, the relevant part of the successor state axiom for B could be replaced by an encoding for dealing with probabilities (as in [Bacchus *et al.*, 1999]). This would potentially yield an approach which combines qualitative and quantitative aspects. Also of interest is the fact that, if the present approach were re-expressed in terms of probabilities, then the separation of the two senses of Alt, as suggested below, would accommodate both subjective and objective probabilities in the same framework.

Further Work Clearly sensors may fail or report an incorrect result. The present approach would seem to extend readily to fallible sensors (and in fact the approach already deals with sensors expected to be correct but which may fail). Intuitively, an incorrect sensor reading can be regarded as a correct reading of some other fluent, perhaps unknown to the user. However, once all actions (physical or sensing) are potentially fallible, the issue arises as to how to deal with actions or classes of actions with differing reliability.

In a quite different direction, it can be noted that in the successor state axiom for B, the two instances of Alt in (2) have distinct senses: the first states what action is actually executed in place of the intended action a_i. Plausibility doesn't play a role here and the plausibility argument plays no role in determining the agent's next belief state. The second instance of Alt is purely epistemic. Plausibility is relevant here and, as the axiom makes clear, we're dealing with subjective plausibilities. Thus the successor state axiom could be generalised, renaming the first occurrence of Alt as $BAlt$ (for *believe Alt*). This new version would allow a more nuanced approach to reasoning about nondeterministic actions, in that one can separate what actions *are* unknowingly executed in place of others, from those actions that the agent *believes* may be executed in place of others. We could then express the case where a dart player believes that he always hits the board, though in fact he is a poor player who is as likely to not hit the board as to hit it.

In this paper we have dealt with a single agent. Given our stance that nondeterminism is an epistemic phenomenon, there is no reason why the approach could not be extended to address actions from other agents, involving perhaps an account of concurrency such as given in [Reiter, 2001].

Finally, we note that there is nothing special about the situation calculus with regards to the approach to nondeterminism. Hence the approach could be readily re-expressed, for example, in terms of an epistemic action language such as [Son and Baral, 2001] where the semantics of the approach is based on a deterministic transition system.

7 Conclusion

This paper has presented a qualitative theory of nondeterminism. The central intuition is that the world is, in fact, deterministic, and that nondeterminism is an artifact of an agent's limited knowledge and perception. The account is based on an epistemic extension to the situation calculus that deals with physical and sensing actions. The account offers several advantages: an agent has categorical (rather than probabilistic) beliefs, yet can deal with equally-likely outcomes (such as in flipping a fair coin), or with outcomes of differing plausibility (such as an action that may on rare occasion fail), or even actions which normally fail and where the desired outcome is the implausible result (such as striking a flint to get a fire). Thus the approach allows for the specification of domain-specific theories of types of action failure, thereby also shedding light on the formal underpinnings of reasoning about action.

Acknowledgements

Yongmei Liu suggested that the Alt predicate be situation dependent. We thank the reviewers for their detailed and pertinent comments. Financial support was gratefully received from the Natural Sciences and Engineering Research Council of Canada.

References

[Bacchus *et al.*, 1999] F. Bacchus, J.Y. Halpern, and H.J. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1-2):171–208, 1999.

[Baral, 1995] Chitta Baral. Reasoning about actions: Non-deterministic effects, constraints, and qualification. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 2017–2026, Montréal, Canada, 1995.

[Boutilier *et al.*, 2000] Craig Boutilier, Raymond Reiter, Mikhail Soutchanski, and Sebastian Thrun. Decision-theoretic, high-level agent programming in the situation calculus. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 355–362, Austin, TX, 2000.

[Cassandra *et al.*, 1994] Anthony R. Cassandra, Leslie Pack Kaelbling, and Michael L. Littman. Acting optimally in partially observable stochastic domains. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 1023–1028, Seattle, WA, 1994.

[Cimatti *et al.*, 2003] Alessandro Cimatti, Marco Pistore, Marco Roveri, and Paolo Traverso. Weak, strong, and strong cyclic planning via symbolic model checking. *Artificial Intelligence*, 147(1-2):35–84, 2003.

[Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[Delgrande and Levesque, 2012] J.P. Delgrande and Hector J. Levesque. Belief revision with sensing and fallible actions. In *Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, Rome, Italy, 2012.

[Gärdenfors, 1988] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.

[Gelfond and Lifschitz, 1998] Michael Gelfond and Vladimir Lifschitz. Action languages. *Electronic Transactions on AI*, 3, 1998.

[Ginsberg and Smith, 1988] Matthew L. Ginsberg and David E. Smith. Reasoning about action ii: The qualification problem. *Artificial Intelligence*, 35(3):311–342, 1988.

[Giunchiglia *et al.*, 2004] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 153(1-2):49–104, 2004.

[Jensen *et al.*, 2004] Rune M. Jensen, Manuela M. Veloso, and Randal E. Bryant. Fault tolerant planning: Toward probabilistic uncertainty models in symbolic nondeterministic planning. In Shlomo Zilberstein, Jana Koehler, and Sven Koenig, editors, *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, pages 335–344. AAAI Press, 2004.

[Laplace, 1814] P. Laplace. *Essai philosophique sur les probabilités*. (English trans: Dover, 1952), Paris, first edition, 1814.

[Levesque *et al.*, 1997] H.J. Levesque, R. Reiter, F. Lin, and R.B. Scherl. Golog: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31, 1997.

[Levesque *et al.*, 1998] H.J. Levesque, F. Pirri, and R. Reiter. Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science*, 3(18), 1998.

[McCarthy, 1977] J. McCarthy. Epistemological problems in artificial intelligence. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1038–1044, Cambridge, MA, 1977.

[Peppas, 2008] P. Peppas. Belief revision. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, San Diego, USA, 2008.

[Poole, 2008] David Poole. The independent choice logic and beyond. In Luc De Raedt, Paolo Frasconi, Kristian Kersting, and Stephen Muggleton, editors, *Probabilistic Inductive Logic Programming*, volume 4911 of *LNAI*, Berlin, 2008. Springer.

[Pratt, 1976] V. Pratt. Semantical considerations on Floyd-Hoare logic. In *17th IEEE Symposium on Foundations of Computer Science*, pages 109–121, 1976.

[Puterman, 1994] M. L. Puterman. *Markov Decision Processes*. Wiley, 1994.

[Reiter, 2001] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, Cambridge, MA, 2001.

[Scherl and Levesque, 2003] R.B. Scherl and H.J. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.

[Shanahan, 1997] Murray Shanahan. *Solving the frame problem - a mathematical investigation of the common sense law of inertia*. MIT Press, 1997.

[Shapiro *et al.*, 2011] Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.

[Son and Baral, 2001] T. Son and C. Baral. Formalizing sensing actions: A transition function based approach. *Artificial Intelligence*, 125(1-2):19–91, 2001.

[Spohn, 1988] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume II, pages 105–134. Kluwer Academic Publishers, 1988.

[Thielscher, 2001] M. Thielscher. The qualification problem: A solution to the problem of anomalous models. *Artificial Intelligence*, 131(1-2):1–37, 2001.