

An Account of Intensional and Extensional Actions, and its Application to Belief, Nondeterministic Actions, and Fallible Sensors

Jens Claßen , James P. Delgrande

School of Computing Science, Simon Fraser University, Burnaby, BC, Canada
jens_classen@sfu.ca, jim@cs.sfu.ca

Abstract

In general, an agent may have incomplete and inaccurate knowledge about its environment. As well, actions may not turn out as intended or may have nondeterministic effects, and sensors may on occasion give incorrect results. We present a general, qualitative approach to reasoning about action and change in such a setting. The approach is expressed as an extension to basic action theories in the situation calculus, where an agent’s epistemic state is modelled by a set of situations, where each situation is assigned a non-negative integer representing its plausibility. The agent’s epistemic state is updated by modifying these plausibility values after the execution of an action, taking into account the possibility of unexpected results. To this end, we consider actions to have an *intensional* aspect, under the control of and determined by the agent, and an *extensional* aspect, not directly accessible to the agent and controlled by “nature”. This leads to two distinct but related notions of belief, an extensional “bird’s eye” view which models an agent’s beliefs wrt actually-executed actions, and an intensional view representing beliefs from the agent’s point of view. We argue that the approach is significantly more general and comprehensive than previous accounts, and leads to a unified view of failed actions and nondeterminism with respect to physical and sensing actions.

1 Introduction

In standard accounts of reasoning about action, an agent has knowledge about the world and it can carry out actions that affect the state of the world. Each action has a *prerequisite* which must hold in order for the action to be executed, and a set of possibly conditional, but fixed and determined, *effects* following the action’s execution. In Knowledge Representation, the best-known and most-studied approach is Reiter’s (2001)¹ account of the *situation calculus*. This approach has been generalised in various ways, notably to include epistemic notions wherein an agent’s *beliefs* evolve following the execution of physical and sensing actions.

These approaches are highly idealised but, under the assumption that “everything works as expected”, are sufficient for formalising notions such as action *progression* and for *planning*. Of course, in general, things rarely work out as expected. An agent’s beliefs about the world may be incomplete, and they may be inaccurate or incorrect. Actions

may fail for a number of reasons: a prerequisite condition may not be satisfied, or an action may simply fail for no known reason. As well, actions may give unexpected or unpredictable results: for example, a gripper may on occasion fail to pick up an object. Moreover, the outcomes of some actions, such as flipping a coin, may be impossible to predict. Similar concerns apply to sensors; for example, a light sensor may fail transiently and unexpectedly.

Our goal in this paper is to formalise the evolution of an agent’s beliefs in such a general, unpredictable, setting. We will also require that the state of the world, including the agent’s beliefs, be well defined following *any* sequence of actions. For example, if an agent attempts to execute a *pickup* action on object o_1 while holding o_2 in its gripper and then go to location l , one would expect that the agent would not be holding o_1 , which would be in its original place, but the agent and o_2 would be at l . Our account will also be *qualitative*, in that an agent will have a set of contingent beliefs that it holds to be true.

In the approach, an agent’s *epistemic state* is modelled by a set of situations (analogous to possible worlds) where each situation has an associated plausibility given by a non-negative integer. Lower-ranked situations are considered more plausible, and the agent’s contingent beliefs are characterised by those situations with minimum plausibility. As actions are carried out, these plausibilities are appropriately modified, resulting in a new epistemic state for the agent. In the most basic case, the approach captures that of (Shapiro *et al.* 2011), which in turn extends the approach of Scherl and Levesque (2003).

A key point of departure of our theory is that an action’s arguments are partitioned into two sets: those that, as usual, the agent controls and can set, and those beyond the control of the agent that are determined by “nature” (we henceforth drop the scare quotes). For example, in flipping a coin, an agent can select a coin and execute a flip action on the coin, but nature decides the outcome of *heads* or *tails*, or perhaps implausibly that the action fails. In general we can express an action a as $a(\vec{x}, \vec{y})$ where \vec{x} is determined by the agent and \vec{y} by nature. For an action $a(\vec{x}, \vec{y})$, we will refer to an action occurrence with just the \vec{x} arguments (controlled by the agent) specified as the *intensional* action, and an occurrence with both \vec{x} and \vec{y} set as the *extensional* action.

Thus, to flip a coin, we could have an action $flip(x, y)$

¹For the most part we defer references to the next section.

where x is a coin and y is the outcome. Hence, y could represent *heads* or *tails* most plausibly, or implausibly *null* for action failure. If *flip* is executed on coin c (the intensional aspect), the agent will not know the outcome, but will believe that it is *heads* or *tails*, while allowing for an implausible failure. The actual outcome of how the coin landed (the extensional aspect) could subsequently be determined by sensing the state of the coin.

We suggest that the approach provides a detailed and nuanced account of actions with potentially unpredictable outcomes, in a qualitative setting. At the most basic, “normal” case, one has a basic action theory in the situation calculus. For example, a *pickup* action with satisfied prerequisites would be expected to yield the expected effects, of the object being picked up, etc. Beyond this, the account allows for the specification of unexpected events: maybe, an object next to the intended object is occasionally picked up instead, or maybe for no known reason the action simply fails. Similarly a sensor may be expected to work correctly, but may sometimes give an incorrect value.

The methodology is conceptually straightforward: An action is axiomatised with regards to what will normally happen along with what may implausibly also happen. Then in executing an action, the agent’s epistemic state is modified to take into account what will most plausibly happen as well as implausibilities. If later evidence shows that its current beliefs cannot be the case, the agent may suitably modify its belief state to take the new evidence into account. This in turn leads to two distinct but related notions of belief, *extensional belief*, which models the agent’s beliefs with respect to actually-executed actions; and *intensional belief*, which gives a subjective account of the agent’s beliefs with regards to what the agent believes it executed.

The next section reviews related work. After this we give an informal description of our approach, followed by the formal development. We then explore properties of the approach. In the final sections we give a general discussion and conclude.

2 Background

The Situation Calculus The situation calculus (Levesque *et al.* 1998; Reiter 2001) is expressed in a first-order language with equality, with sorts for actions, situations, and objects (i.e. everything else). Here we review the epistemic extension to the situation calculus of Scherl and Levesque (2003). A *situation* represents a finite world history, composed of a sequence of actions from some initial state of the world. Predicates and functions whose values are situation-dependent are called *fluents*; by convention the last argument of a fluent is a situation. For example, $\neg \text{Held}(x, s)$ might assert that the agent is not holding object x in situation s . A set of initial situations expresses the ways the domain might be initially according to the agent, while the constant S_0 denotes the actual initial state of the world. The term $do(a, s)$ denotes the unique situation resulting from executing action a in situation s ; for example $\text{Held}(x, do(\text{pickup}(x), s))$ asserts that the agent is holding object x following a *pickup* action in s .

A dynamic domain is axiomatised in terms of a *basic action theory* (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; (3) successor state axioms (SSA), 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.

An agent’s knowledge is axiomatised by considering situations as possible worlds. Two special fluents are 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.² The B fluent is the usual belief accessibility relation of modal logic: $B(s', s)$ holds when the agent in situation s thinks that situation s' might be the actual situation.³ A successor state axiom for B is given by:⁴

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

Thus 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 is defined, as usual, as truth in all accessible situations:⁵

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

Belief Change and the Situation Calculus In the AGM approach to belief revision (Gärdenfors 1988; Peppas 2008) an agent’s epistemic state is modelled by a total preorder over possible worlds, where the agent’s (contingent) belief set is characterised by the minimally-ranked worlds. In revising by a formula ϕ , the agent’s new belief set is characterised by the minimal ϕ worlds; different recipes have been proposed for how a new total preorder is obtained; see (Darwiche and Pearl 1997; Peppas 2008). A somewhat more general scheme is to assign non-negative integers to possible worlds, where an integer gives the plausibility of a world, with lower-ranked worlds being more plausible (Spohn 1988).

This latter scheme is adopted in (Shapiro *et al.* 2011) in the situation calculus, where situations are assigned a plausibility. Their approach treats actions as in (Scherl and Levesque 2003), and it requires sensing to be exact. Delgrande and Levesque (2012; 2019) treat sensing like belief revision. In addition they allow for an agent to inadvertently execute an unintended action. This results in a 4-place B

²Actions that return no useful sensing information like *pickup(c)* are axiomatised as always returning true.

³Because B is a fluent, the order of the situation arguments is reversed from the usual convention in modal logic.

⁴Here and elsewhere, free variables are assumed to be universally quantified from the outside.

⁵ 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 .

fluent with a (in their words) daunting successor state axiom. (Klassen *et al.* 2018; Klassen *et al.* 2020) also address belief change in the situation calculus, but plausibility levels there are defined by means of cardinality-based circumscription; essentially abnormalities are counted and situations with the fewest abnormalities are considered most plausible. Among other work, (Bacchus *et al.* 1999) addresses noisy sensors and actuators in a probabilistic framework, while (Schwering *et al.* 2017) deal with conditional beliefs, covering natural revision and lexicographic revision in the epistemic situation calculus (Lakemeyer and Levesque 2004); as with (Shapiro *et al.* 2011) they assume that sensing is correct. Fang and Liu (2013) address belief change in a multiagent setting in the situation calculus, while Fang, Liu and Wen (2015) study the progression of belief for nondeterministic actions in propositional situation calculus theories. Shapiro and Pagnucco (2004) address the possible discrepancy between belief and sensing by allowing the agent to hypothesise exogenous actions to account for any difference.

A Limitation Concerning Knowledge In the standard approach to the situation calculus, the *result* of executing an action (given in the effect axioms) is treated separately from the possibility (or *legality*) of that action. While this may work well for planning, it is inadequate for dealing with incomplete or incorrect beliefs. Consider the following example: there is a robot arm with an attached electromagnet that the agent cannot observe. A *pickup* action of an object is possible if the arm is down, the current is on, and the object is made of iron and is nearby. Assume that the arm is down, there is an iron object nearby, and the current is on. The agent knows the first two items, but it does not know whether the current is on. However, after executing an unobserved *pickup* action, the agent will believe that the current is on and the object has been lifted. This is clearly unwarranted.

The problem is that in the successor state axiom (1) for B , those situations that are not possible wrt an executed action are “pruned” from the B accessibility relation. This difficulty⁶ isn’t resolved by removing $Poss(a, s^*)$ from (1) since in this case, in all successor situations, possible or not, the effects of *pickup* will hold. To the best of our knowledge, all accounts of belief in the situation calculus are limited, in that in dealing with incomplete information an agent may draw conclusions that are unwarranted. This includes, among others (Levesque 1996; Levesque *et al.* 1998; Bacchus *et al.* 1999; Reiter 2001; Scherl and Levesque 2003; Fang and Liu 2013; Rajaratnam *et al.* 2014). In addition, (Shapiro *et al.* 2011; Delgrande and Levesque 2012; Schwering *et al.* 2017; Delgrande and Levesque 2019) avoid the problem by assuming that all actions are possible in all situations.

3 The Approach: Intuitions

Our goal is to model the evolution of an agent’s epistemic state, where an epistemic state is characterised by a set of

⁶The full paper will present a detailed argument; as well, a supplement with the details is available from the authors.

situations with associated plausibility values. To begin, we adopt Reiter’s (2001) account of the situation calculus but with a modest modification: For the effect axioms for a fluent F with respect to action a , we include the precondition of action a in the antecedent of the effect axioms. This has the result that all action outcomes are well-specified in all circumstances: if an action’s preconditions are not satisfied, then on *do*-ing the action, the fluent in question is unchanged, i.e. “nothing happens”. As argued earlier, this is required for an adequate account of belief and action in the case of incomplete knowledge.

An agent’s epistemic state is modelled by a set of situations where each situation is labelled with a non-negative integer indicating its (epistemic) plausibility. For this we use a fluent $B(s', p, s)$ that asserts that in situation s the agent attaches plausibility p to s' . The agent’s contingent beliefs at s are characterised by those B -related situations s' with minimum plausibility.⁷ The key challenge is to specify what happens to this ranking (and so the agent’s epistemic state) as physical and sensing actions are executed. We note that in addressing this question, we remain within the paradigm of reasoning about action in the situation calculus; specifically we do not import notions nor desiderata from belief revision (beyond having a plausibility ranking over situations). See Section 6 for a discussion.

Actions are deterministic, as is stipulated in the successor-state axioms which specify actions’ effects on fluents. However, as described, our goal is to allow for (apparent) non-determinism, whether in an action that may unexpectedly fail or give abnormal results, or in executing an action with unpredictable, seemingly nondeterministic results. To this end, our stance is that the world is deterministic, but due to its limited capabilities, an agent may not know the full consequences of its actions. For example, in flipping a coin, an agent will select a coin and in the most plausible case execute a “flip heads” or a “flip tails” action. However, due to imprecise motor control and limited sensing, the agent will not know the result of a flip until observing its result.

To this end, we develop a methodology in which the agent, intending to execute a particular action, has control over some of the action’s arguments, while other arguments are controlled by nature. The former set of arguments are just those that would appear in a standard situation calculus theory. The latter are beyond the control of the agent, and may introduce uncertainty or noise in the effects of that action. As described, this can be regarded as nature interfering with what the agent intends by determining the value of an action argument. For example, if the agent intends to push a switch at position i , nature may interfere so that with plausibility $|i - j|$, a switch at j is pushed. We implement this by considering action arguments as being of one of two disjoint types: those that the agent can control, and those that only nature can affect. We will refer to an action along with values that the agent can control as an *intensional action*, while an action in which all arguments have values as an *exten-*

⁷Most approaches ensure that some B -accessible situation has 0 plausibility. For us it is more convenient to allow the minimum plausibility to be non-zero; see Example 1 in the next section.

sional action. Thus the action of the agent flipping a given coin is an intensional action, whereas the action of flipping a given coin yielding heads is an extensional action. In general then, for an action a , the intensional action corresponding to a is that part or aspect of the action that the agent has control over; this is the action that the agent would identify as having executed. The extensional action corresponding to a would be the action that was in fact executed. In an action theory, the intensional and extensional aspects are related by having some extensional aspects be *expected* or *most plausible*. Thus, in the large majority of cases, actions will have their expected results – a pickup action will succeed, a sensor will be correct, etc. Note however that there is nothing to prevent the encoding of an action that is expected to fail. For example, in striking a flint it may be that normally nothing happens; however, implausibly, it may be that a fire is started with a flint strike.

To encode these possibilities we introduce a predicate $Apl(a(\vec{x}, \vec{y}), p, s)$. Apl is mnemonic for “(action) argument plausibility”; a is an action with arguments \vec{x} set by the agent and \vec{y} set by nature. The intended interpretation of $Apl(a(\vec{x}, \vec{y}), p, s)$ is that for specific values of \vec{x} set by the agent, specific values of \vec{y} have plausibility p in situation s . For example, in flipping a coin where x is a coin and y the result, we might have:

$$Apl(flip(x, y), p, s) \equiv ((y = h \vee y = t) \wedge p = 0) \vee (y = null \wedge p = 1) \quad (2)$$

Thus, with plausibility 0 the result will be heads or tails, and with plausibility 1 the action will fail.

Instances of Apl are used in the successor-state axiom of the B fluent. We differ from most approaches to belief change, in that all alternative outcomes of an action must be explicitly encoded. An example where this is not the case is in classical belief revision: on revising by ϕ , the $\neg\phi$ worlds are re-ranked so that there is no minimal $\neg\phi$ world. In our approach, if a sensor (correctly) reports that ϕ is the case, then it must be that ϕ holds at all B -related successor situations of this action. This is what is done in (Shapiro *et al.* 2011), where sensing is always correct. However, we generalise (Shapiro *et al.* 2011) in that one can also axiomatise that sensing may implausibly fail in some fashion, or return an incorrect result.

The overall result is that actions are deterministic and we remain within Reiter’s conception of the situation calculus. However, from the agent’s point of view the domain may appear nondeterministic or result in unexpected outcomes. This also means that the agent’s epistemic state is well-defined in any situation: For situation s , the agent’s epistemic state is given by the entailed instances of $B(s', n, s)$, and the agent’s belief set is characterised by those s' for which n is minimal. Thus any situation s can be thought of as indexing an epistemic state (the entailed instances of the form $B(s', n, s)$), and as the world evolves, the agent’s epistemic state evolves along with it. We refer to this conception of belief as *extensional belief*, since it corresponds to what an external, omniscient observer would observe as actions are successively executed. It is also, of course, the standard conception of belief and belief evolution in epis-

temic versions of the situation calculus.

However, one can also consider the evolution of an agent’s beliefs solely from the point of view of the agent, and with no reference to any external “actual” situation. For example, consider an agent with no contingent knowledge of the domain, and with a light sensor that usually works correctly, but occasionally reports *on* or *off* independent of the actual situation. If the agent performs a sensing action and the result is *off*, the agent will believe that the light is off, since this is the most plausible outcome. If two subsequent sensing actions report *on*, the agent may come to believe that the light is in fact on, and that the first action failed. We refer to this notion as *intensional belief*, since the agent has no direct access to an underlying “actual” situation, only the reported results of its (fallible) sensors and actuators.

4 The Approach: Formal Details

4.1 Basic Action Theories

We begin with a *basic action theory* (BAT), as summarised in Section 2, but with one adjustment: We assume that an action’s preconditions are included in the successor state axiom for a fluent with respect to that action. The easiest way that this can be enforced is to assume for action $a(\vec{x})$ and fluent F that there is a positive effect axiom of the form

$$Poss(a(\vec{x}), s) \wedge \varepsilon_F^+ \supset F(\vec{y}, do(a(\vec{x}), s))$$

and an analogous negative effect axiom.⁸ Then Reiter’s (1991; 2001) solution to the frame problem “compiles” the effect axioms into the successor state axiom. This has the result here that if an action’s preconditions are not satisfied, and the agent executes the action, nothing happens.

As described, nature may influence an action by determining some of its arguments. *How* this is done is specified in the theory axiomatisation. These arguments are not directly accessible to the agent but may potentially be discovered by sensing. A characteristic of such arguments is that they play no role in that action’s precondition.⁹ Thus for action a , with \vec{x} being regular action arguments and \vec{y} being those controlled by nature, an action precondition is a sentence of the form

$$Poss(a(\vec{x}, \vec{y}), s) \equiv \Pi_a(\vec{x}, s)$$

where $\Pi_a(\vec{x}, s)$ is a formula that is *uniform in s* (i.e., it does not mention any situation terms apart from s), and whose free variables are among \vec{x}, s .

We illustrate these notions in an example that deals with flipping a coin. Assume that we have the predicate $Coin(x)$, that object x is a coin, and the fluents: $NextTo(x, s)$: the agent is next to object x in situation s ; $Held(x, s)$: object x is held by the agent in situation s ; $Shows(x, z, s)$: coin x shows the result z in situation s where $z \in \{h, t, s, null\}$,

⁸As noted, most often action preconditions are kept distinct from effect axioms, e.g. (Reiter 2001; Brachman and Levesque 2004), but not always (Reiter 1991).

⁹They may however depend on the state of the environment. For example an agent may be accurate in throwing an object in good light, but inaccurate in poor light.

standing for “heads”, “tails”, “on its side”, and “null” respectively (and where “null” would apply to a coin that remains held). The *flip* action will have two arguments; *flip*(x, y) will be the action of a flip of coin x , where the outcome y is controlled by nature and $y \in \{h, t, s, null\}$. A flip action is deterministic, and *flip*(c, h) is the action of flipping coin c such that it shows heads. The idea, developed below, is that an agent will execute a deterministic action like “flip heads”, but it will not know whether it executed that action or a “flip tails” action. The former is what is referred to as an *intensional action* while the latter is an *extensional action*.

The *flip* action has precondition: $Poss(\text{flip}(x, y), s) \equiv Coin(x) \wedge Held(x, s)$. Note that the argument controlled by nature, y , does not appear on the right hand side of the equivalence. To complete the example there are two further actions: *drop*, whatever is *Held* is dropped, and with precondition $Poss(\text{drop}, s) \equiv \exists y. Held(y, s)$; and *pickup*(x) where object x is to be picked up, and with precondition $Poss(\text{pickup}(x), s) \equiv NextTo(x, s) \wedge \neg \exists z. Held(z, s)$.

Let $z \in \{h, t, s\}$ stand for ($z = h \vee z = t \vee z = s$). We have the effect axioms:

$$Held(x, s) \wedge Coin(x) \wedge z \in \{h, t, s\} \supset$$

$$Shows(x, z, do(\text{flip}(x, z), s))$$

$$NextTo(x, s) \wedge \neg \exists y. Held(y, s) \supset$$

$$Held(x, do(\text{pickup}(x), s))$$

$$Held(x, s) \wedge Coin(x) \wedge z \in \{h, t, s\} \supset$$

$$\neg Held(x, do(\text{flip}(x, z), s))$$

$$\exists y. Held(y, s) \supset \neg Held(x, do(\text{drop}, s))$$

From these we can derive the SSAs (with simplification):

$$Shows(x, y, do(a, s)) \equiv [Held(x, s) \wedge Coin(x) \wedge a = \text{flip}(x, y) \wedge y \in \{h, t, s\}] \vee Shows(x, y, s)$$

$$Held(x, do(a, s)) \equiv [NextTo(x, s) \wedge \neg \exists y. Held(y, s) \wedge a = \text{pickup}(x)] \vee [Held(x, s) \wedge \neg (\exists z (a = \text{flip}(x, z) \wedge Coin(x) \wedge z \in \{h, t, s\}) \vee a = \text{drop}(x))]$$

From the successor state axiom of *Shows* we can derive:

$$(\neg Held(x, s) \vee a = \text{flip}(x, null)) \supset$$

$$(Shows(x, z, s) \equiv Shows(x, z, do(a, s)))$$

Hence nothing happens in case of action failure, as expected.

4.2 Sensing

Our approach to sensing is similar to that of (De Giacomo and Levesque 1999), using *guarded fluent axioms*. There, $SF(a, s)$ holds just if the sensor associated with a is true in situation s . For example, take $L(x, s)$ to hold iff the light is on in location x in situation s . Then we might use the action *sLt* to sense whether the light is on in the room where the agent is located. We could then add an argument controlled by nature, with possible values *ok*, *t.off*, and *t.on*, where the last two erroneously and transiently report the light being off or on, respectively. Then we could have:

$$At(x, s) \supset (SF(\text{sLt}(\text{ok}), s) \equiv L(x, s)) \quad (3)$$

$$\neg SF(\text{sLt}(\text{t.off}), s) \quad (4)$$

$$SF(\text{sLt}(\text{t.on}), s) \quad (5)$$

More generally, for each sensor we have a set of axioms where there is one axiom for each choice of arguments made by nature. A sensor axiom is of the following form, where γ is a guard, and a a sensing action for fluent F .

$$\gamma(\vec{x}, s) \supset (SF(a(\vec{x}, \vec{y}), s) \equiv \{\neg\}F(\vec{x}, \vec{y}, s)) \quad (6)$$

Strictly speaking (4) would then be written something like $\top \supset (SF(\text{sLt}(\text{t.off}), s) \equiv \perp)$. A light sensor that gives a parity error could be encoded as $At(x, s) \supset (SF(\text{sLt}(\text{parity.err}), s) \equiv \neg L(x, s))$. Last, for completeness, if k ranges over axioms (6) with guard γ_k , we have:

$$Poss(a(\vec{x}, \vec{y}), s) \equiv \bigvee_k \gamma_k(\vec{x}, s).$$

4.3 Action Plausibility

We add two predicates to the theory to deal with action plausibility: $Apl(a, p, s)$ (for (action) *argument plausibility*) specifies the plausibility p attached to specific argument values of action a determined by nature at situation s , and $Ieq(a, a')$ asserts that a and a' are intensionally equivalent, i.e. the arguments set by the agent are the same in a and a' .

We deal with *Ieq* first, since it is simpler. We have the axioms:

$$\neg Ieq(a(\vec{x}), a'(\vec{y})) \quad \text{for } a, a' \text{ distinct function symbols}$$

$$Ieq(a(\vec{x}_1, \vec{y}_1), a(\vec{x}_2, \vec{y}_2)) \equiv \vec{x}_1 = \vec{x}_2$$

where for action $a(\vec{x}, \vec{y})$, \vec{x} is set by the agent and \vec{y} is controlled by nature.

Apl is a special predicate (like *Poss*). For an action, axioms of the following form are provided, where as usual for $a(\vec{x}, \vec{y})$, a is an action function symbol, \vec{x} is determined by the agent, and \vec{y} by nature; the formula Ψ_a is uniform in s .

$$Poss(a(\vec{x}, \vec{y}), s) \supset (Apl(a(\vec{x}, \vec{y}), p, s) \equiv \Psi_a(\vec{x}, \vec{y}, p, s))$$

$$\neg Poss(a(\vec{x}, \vec{y}), s) \supset (Apl(a(\vec{x}, \vec{y}), p, s) \equiv p = 0)$$

The second part is included for completeness; if an action’s preconditions are not satisfied, then it cannot be executed. In this case (as discussed in Section 4.1), the action has no effect, and there is only one possible outcome in which all fluents remain unchanged. For example a *flip* action cannot be (successfully) executed on coin c if c is not held. In such a case, the *flip* action would have no effect; this is the only possible outcome and the plausibility in this case is $p = 0$.

Clearly plausibility is a function of the action in *Apl*. We have the axiom:

$$Apl(a, p_1, s) \wedge Apl(a, p_2, s) \supset p_1 = p_2.$$

We assume for any action theory with action instance $a(\vec{t}, \vec{y})$ and situation s , that there is an instance $a(\vec{t}, \vec{t}')$ such that $Apl(a(\vec{t}, \vec{t}'), 0, s)$ is entailed. That is, at every situation, every intensional action instance has an extensional instance with plausibility 0, meaning *something* (perhaps nothing) is expected to happen.

Examples: Let us return to our earlier example of flipping a coin. We can assert that normally flipping a coin results in heads or tails, although implausibly the action may fail, and even more implausibly the coin will land on its side:

$$\begin{aligned} Poss(flip(x, y), s) \supset [Apl(flip(x, y), p, s) \equiv \\ ((y = h \vee y = t) \wedge p = 0) \vee \\ (y = null \wedge p = 1) \vee (y = s \wedge p = 2)] \end{aligned}$$

For pushing a light switch on a bank of switches, where x is the switch that the agent intends to push and y is the actual switch, determined by nature:

$$\begin{aligned} Poss(push(x, y), s) \supset \\ [Apl(push(x, y), p, s) \equiv p = |x - y|] \end{aligned} \quad (7)$$

Throwing a dart at a board where the argument x is controlled by nature. If it's not dim, the dart hits the board; if it is dim, the dart is as likely to hit the wall as the board.

$$\begin{aligned} Poss(throw(x), s) \supset \\ Apl(throw(x), p, s) \equiv (x = board \wedge p = 0) \wedge \\ (\neg Dim(s) \supset (x = wall \wedge p = 1)) \wedge \\ (Dim(s) \supset (x = wall \wedge p = 0)) \end{aligned}$$

We can use Apl to take care of normally-functioning sensors along with transient errors. Consider our earlier light-sensing example (3–5). We have the following.

$$\begin{aligned} Poss(sLt(x), s) \supset [Apl(sLt(x), p, s) \equiv [((x = ok \wedge \\ p = 0) \vee ((x = t.off \vee x = t.on) \wedge p = 1))] \end{aligned}$$

Thus, most plausibly ($p = 0$) the sensor will report correctly; implausibly ($p = 1$) there will be a transient error.

Similarly, in analogy to pushing the wrong switch in a bank of switches (7), we can encode sensing a position incorrectly. Below, x is the position sensed by the agent and y is the true position.

$$\begin{aligned} SF(sPos(x, y), s) \equiv y = pos(s) \\ Apl(sPos(x, y), p, s) \equiv p = |x - y| \end{aligned}$$

4.4 Extensional Belief

Belief is expressed by a 3-place version of an accessibility relation, where $B(s', p, s)$ has the interpretation that, in situation s , as far as the agent is concerned, situation s' could be the actual situation with plausibility p . We assume that if an agent is in an initial situation, then all accessible situations are also initial. As well, we stipulate that the initial plausibility value is a function of a situation pair. (It will subsequently be shown that this functional relation will hold for all B -accessible situations.) We have the axioms:

$$\begin{aligned} Init(s) \wedge B(s', n, s) \supset Init(s') \\ Init(s) \wedge B(s', n_1, s) \wedge B(s', n_2, s) \supset n_1 = n_2 \end{aligned}$$

There are various possibilities for a successor state axiom for B , and the actual SS axiom will reflect the KB designer's

revision strategy.¹⁰ Perhaps the most basic is the following.

$$\begin{aligned} B(s', n, do(a, s)) \equiv \exists p', s^*, n^*, a^*, p^*. B(s^*, n^*, s) \wedge \\ Apl(a, p', s) \wedge Ieq(a, a^*) \wedge Apl(a^*, p^*, s^*) \wedge s' = do(a^*, s^*) \\ \wedge (SF(a^*, s^*) \equiv SF(a, s)) \wedge n = n^* + p' + p^* \end{aligned} \quad (8)$$

The equivalence spells out when situation s' has plausibility n following the execution of action a in s . The set of all such s' determines the agent's epistemic state at $do(a, s)$. We have the following. First, the B instance on the left hand side is obtained from an instance, $B(s^*, n^*, s)$. The action a has plausibility p' ; that is, the argument values set by nature have plausibility p' . The second occurrence of Apl along with Ieq asserts that, for all the agent knows a^* was executed where the arguments set by nature have plausibility p^* . Then s' is a s^* successor given by action a^* . The overall plausibility n is given by the original plausibility n^* , the plausibility p' of the executed action a , and the plausibility p^* of a^* . As well, the respective situations s and s^* must agree on the result of the sensing action.

Given this, we define that a formula ϕ is believed just if it holds at all accessible situations with minimum plausibility.

$$\begin{aligned} MP(s', s) \doteq \exists p. B(s', p, s) \wedge \forall s'', p'(B(s'', p', s) \supset p \leq p') \\ Bel(\phi, s) \doteq \forall s'. MP(s', s) \supset \phi[s']. \end{aligned}$$

Example 1. There are two situations, where the light is on in one but not the other: $L(S_0)$, $\neg L(S_1)$. The agent has no knowledge of the state of the light: $B(S_0, 0, S_0)$, $B(S_1, 0, S_0)$. If the agent senses the light, then with plausibility 0 the sensing will be correct, that is $sLt(ok)$ is executed. For $B(s', n, do(sLt(ok), S_0))$, there are six candidate instances; for $s' = do(a^*, s^*)$ we may have $a^* \in \{sLt(ok), sLt(t.off), sLt(t.on)\}$, and $s^* \in \{S_0, S_1\}$. These six instances must also agree with $SF(sLt(ok), S_0)$; three of them do, and it can be verified that the following instances are entailed:

$$\begin{aligned} B(do(sLt(ok), S_0), 0, do(sLt(ok), S_0)) \\ B(do(sLt(t.on), S_0), 1, do(sLt(ok), S_0)) \\ B(do(sLt(t.on), S_1), 1, do(sLt(ok), S_0)) \end{aligned}$$

Hence, the agent believes (correctly) that the light is on, and implausibly that the light could be reported on by a transitory on error.

Next, assume that the agent senses the light, but implausibly ($p' = 1$) an off error occurs. For $B(s', n, do(sLt(t.off), S_0))$, there are six instances where $s' = do(a^*, s^*)$, $a^* \in \{sLt(ok), sLt(t.off), sLt(t.on)\}$, and $s^* \in \{S_0, S_1\}$. These six instances must also agree with $SF(sLt(t.off), S_0)$. Once again, three of them do, and it can be verified that the following instances are entailed:

$$\begin{aligned} B(do(sLt(ok), S_1), 1, do(sLt(t.off), S_0)) \\ B(do(sLt(t.off), S_0), 2, do(sLt(t.off), S_0)) \\ B(do(sLt(t.off), S_1), 2, do(sLt(t.off), S_0)) \end{aligned}$$

¹⁰For example, in (8) we choose to not necessarily satisfy the AGM success postulate. It is straightforward to encode AGM-style revision; see for example (Delgrande and Levesque 2012; Delgrande and Levesque 2019) for encodings of approaches of (Nayak et al. 2003) and (Darwiche and Pearl 1997) respectively.

The agent believes (incorrectly) that the sensor reported the correct result. Hence the agent believes the light is off, and implausibly that the light could be reported as off by a transitory off error. Note that the accessible situation with minimal plausibility (viz. $do(sLt(ok), S_1)$) has plausibility 1, reflecting the abnormal action.

This formulation of belief is called *extensional* since it provides an “external” view of the agent’s belief state taking into account the actual actions executed. This view assumes that there is a situation corresponding to the real world, with which sensing is with respect to. Thus, if an agent correctly sensed the state of a light, then it would believe that the light is on; if the sensor gave a transitory *off* reading, then the agent would (erroneously) believe that the light is off. In both cases, following a sensing action the agent believes that the action was successfully carried out.

Progression, that is, determining what the agent would believe given that a sequence of (extensional) actions are executed is conceptually straightforward: From the axiomatisation one can determine what each situation would look like and what situations would be *B*-accessible following any (extensional) sequence of actions.

4.5 Intensional Belief

There is another, arguably more useful and intuitive notion of “projection” that involves considering actions executed from the agent’s point of view. As an example, consider the scenario of Example 1, but where the agent senses the state of the light three times, and where the sensor reports *off*, *on*, and *on*. Clearly, not all of the sensing actions are correct, and the simplest explanation is that the first sensing action had a transitory *off* failure. It thus seems reasonable that after these actions the agent would believe that the light is on. Note that the “straightforward” projection of actions for extensional belief above won’t work here since, in effect, it is only after subsequent sensing actions that the agent would come to believe that the first action had failed. So here we consider the case where the agent executes a sequence of intensional actions and receives sensor reports, and the goal is to determine the resulting epistemic state.

To do this we adjust our background BAT in two ways. The first adjustment addresses the subjective conception of sensing in this approach. In the usual case, an agent carries out a sensing action and the result is determined with respect to the underlying “actual” situation. Here in contrast, we deal with the *report* of a sensor. That is, sensing will be of the form that a sensor reports *that* a fluent is true or *that* a fluent is false. Importantly, the “actual situation” doesn’t play a role; we’re simply dealing with the report of a sensor. To deal with this, for each sensing action, such as $sLt(x)$, we introduce two sensing report actions which will be named by appending a *T* or *F* to that sensing action name. For example we will introduce $sLtT(x)$ and $sLtF(x)$. The interpretation will be that $sLtT(x)$ is the action in which the sensor reports that the light is on, and $sLtF(x)$ is the sensing report that the light is off. We introduce the special fluent $ST(a, s)$, analogous to $SF(a, s)$, to express *that* the sensor report given by *a* holds in *s*. Then we specify the effect of

these actions as follows:

$$\begin{aligned} ST(sLtT(x), s) &\equiv SF(sLt(x), s) && \text{and} \\ ST(sLtF(x), s) &\equiv \neg SF(sLt(x), s) \end{aligned} \quad (9)$$

Since this looks to be rather trivial, some explanation is in order. The use of $sLtT(\cdot)$ is indeed unnecessary, but we choose to use it, first as a reminder that we are encoding a sensor report, and second as a dual to $sLtF(\cdot)$. The action $sLtF(\cdot)$ is more meaningful since, as a sensor result, it will appear in the *ST* fluent: an expression like $ST(\neg sLt(x))$ is ill-formed, whereas $ST(sLtF(x))$ is not.

The second modification concerns the *B* fluent. For extensional belief, this fluent was expressed as $B(s', n, s)$. Each *s* in a fluent instance indexes an agent’s belief state, implicitly given by the various “accessible” *s'* and associated plausibilities *n*. Here we index an agent’s epistemic state not by a situation, but by the (intensional) actions that the agent executed. For a regular action, this will consist of the intensional part of the action; that is for an action $a(\vec{x}, \vec{y})$ with \vec{y} controlled by nature, we will be concerned with the intensional part given by $a(\vec{x}, \cdot)$. As well, we will consider sensing reports, as described above.

The resulting predicate is denoted B^i and is of the form $B^i(s, n, l)$. The first two arguments are as with *B*. The last argument is a list of intensional actions, where a list is specified as:

1. ϵ : the empty list.
2. $ido(a, l)$ where *l* is a list and *a* is an intensional action. This will be either a regular action where only the arguments set by an agent are meaningful, or a sensing report.

We assume that the action theory includes a set of foundational axioms for *ido* analogous to those for *do*. The agent’s initial epistemic state is given by a set of instances of the form $B^i(s, n, \epsilon)$, where *n* is a function of *s*. This gives the axioms:

$$\begin{aligned} B^i(s', n, \epsilon) &\supset Init(s') \\ B^i(s', n_1, \epsilon) \wedge B^i(s', n_2, \epsilon) &\supset n_1 = n_2 \end{aligned}$$

There are again various possibilities for the successor state axiom for B^i , reflecting the KB designer’s revision strategy. Perhaps the most basic is the following.

$$\begin{aligned} B^i(s', n, ido(a, l)) &\equiv \exists s^*, n^*, a^*, p^*. B^i(s^*, n^*, l) \wedge \\ &Ieq(a, a^*) \wedge Apl(a^*, p^*, s^*) \wedge s' = do(a^*, s^*) \wedge \\ &ST(a^*, s^*) \wedge n = n^* + p^* \end{aligned} \quad (10)$$

Thus, for the execution of an action *a* regarded intensionally, one considers the ways that the action may be realised extensionally via *Ieq* and *Apl*. For such an extensional realisation a^* with plausibility p^* , there is a B^i -related situation given by $s' = do(a^*, s^*)$ with overall plausibility given by that of the underlying B^i instance (viz. n^*) and that of the action a^* (viz. p^*). For a sensing report, the situation in question s^* must agree with the report, as given by $ST(a^*, s^*)$.

Example 2. Consider the Example 1 scenario with two situations with the light on in one but not the other, $L(S_0)$, $\neg L(S_1)$, and where the agent does not know the state of

the light: $B^i(S_0, 0, \epsilon)$, $B^i(S_1, 0, \epsilon)$. The agent performs 3 light-sensing actions, with reports off, on, on. Following the action $ido(sLtF(x), \epsilon)$ (i.e. the sensor reports the light is off), the following instances are entailed:

$$\begin{aligned} B^i(do(sLtF(ok), S_1), 0, ido(sLtF(x), \epsilon)) \\ B^i(do(sLtF(t.off), S_0), 1, ido(sLtF(x), \epsilon)) \\ B^i(do(sLtF(t.off), S_1), 1, ido(sLtF(x), \epsilon)) \end{aligned}$$

The agent believes that the light is off, but (implausibly) that the off report is due to a sensor error. The agent again senses the light, and the sensor reports that the light is on. Let $l_1 = ido(sLtT(x), ido(sLtF(x), \epsilon))$. It can be verified that the following instances are entailed (where for better readability we write $do(a_2, do(a_1, s))$ as $do(\langle a_1, a_2 \rangle, s)$):

$$\begin{aligned} B^i(do(\langle sLtF(t.off), sLtT(ok) \rangle, S_0), 1, l_1) \\ B^i(do(\langle sLtF(t.off), sLtT(t.on) \rangle, S_0), 2, l_1) \\ B^i(do(\langle sLtF(ok), sLtT(t.on) \rangle, S_1), 1, l_1) \\ B^i(do(\langle sLtF(t.off), sLtT(t.on) \rangle, S_1), 2, l_1) \end{aligned}$$

The agent is now agnostic about the state of the world, as given by the two minimal (with plausibility 1) situations. Once again the agent senses the light, and the sensor reports that the light is on. Let $l_2 = ido(sLtT(x), ido(sLtT(x), ido(sLtF(x), \epsilon)))$. It can be verified that the following instances are entailed (abbreviating $do(a_3, do(a_2, do(a_1, S_0)))$ as $do(\langle a_1, a_2, a_3 \rangle, S_0)$):

$$\begin{aligned} B^i(do(\langle sLtF(t.off), sLtT(ok), sLtT(ok) \rangle, S_0), 1, l_2) \\ B^i(do(\langle sLtF(t.off), sLtT(ok), sLtT(t.on) \rangle, S_0), 2, l_2) \\ B^i(do(\langle sLtF(t.off), sLtT(t.on), sLtT(ok) \rangle, S_0), 2, l_2) \\ B^i(do(\langle sLtF(t.off), sLtT(t.on), sLtT(t.on) \rangle, S_0), 3, l_2) \\ B^i(do(\langle sLtF(ok), sLtT(t.on), sLtT(t.on) \rangle, S_1), 2, l_2) \\ B^i(do(\langle sLtF(t.off), sLtT(t.on), sLtT(t.on) \rangle, S_1), 3, l_2) \end{aligned}$$

Consequently, the agent believes that the light is on, and the first sensing action failed.

5 Properties of the Framework

In this section, Σ will denote a basic action theory containing the axioms above and using the abbreviations introduced.

The first property is the minor result, claimed earlier, that for extensional belief, plausibility is a function of the situation at hand and an accessible situation. A similar result holds for intensional belief.

Theorem 1.

$$\begin{aligned} \Sigma \models (B(s', n_1, s) \wedge B(s', n_2, s)) \supset n_1 = n_2 \\ \Sigma \models (B^i(s', n_1, l) \wedge B^i(s', n_2, l)) \supset n_1 = n_2 \end{aligned}$$

Proof. By induction over the length of s or l resp. \square

The next property relates B and B^i . For this purpose, assume that Σ satisfies the following additional assumptions. First, B and B^i agree initially:

$$Init(s) \supset [B(s, n, S_0) \equiv B^i(s, n, \epsilon)] \quad (11)$$

Next, “sensing-that” actions (such as $sLtT(x)$, $sLtF(x)$) are compatible with “sensing-if” actions (such as $sLt(x)$) as follows. There are functions¹¹ $o(\cdot)$ and $r(\cdot)$ such that for any sensing-that action a , $o(a)$ extracts the corresponding sensing-if action (e.g., $o(sLtF(x)) = sLt(x)$), and $r(a)$ extracts the sensing result (e.g., $r(sLtF(x)) = \perp$). For any combination of sensing-if action and sensing result, a corresponding sensing-that action exists:

$$\forall a, s \exists a'. a = o(a') \wedge SF(a, s) \equiv r(a') \quad (12)$$

A binary relation $Map(s', s)$ that relates situations s' consisting of sensing-if actions to situations s consisting of the corresponding sensing-that actions is given through the following initial axiom and SSA:

$$Init(s) \supset [Map(s', s) \equiv s' = s] \quad (13)$$

$$Map(s', do(a, s)) \equiv \quad (14)$$

$$\exists a', s''. s' = do(a', s'') \wedge Map(s'', s) \wedge a' = o(a)$$

Compatible actions then agree on action plausibility:

$$Map(s', s) \supset [Apl(a, p, s) \equiv Apl(o(a), p, s')] \quad (15)$$

ST is defined in terms of SF as follows (similar to (9)):

$$Map(s', s) \supset [ST(a, s) \equiv (SF(o(a), s') \equiv r(a))] \quad (16)$$

Intensional equivalence for sensing-that actions corresponds to equivalence of the corresponding sensing-if actions and sensing results:

$$Ieq(a_1, a_2) \equiv Ieq(o(a_1), o(a_2)) \wedge (r(a_1) \equiv r(a_2)) \quad (17)$$

We next define a 3-place predicate $C(l, p, s)$ that serves a similar purpose as $Map(\cdot, \cdot)$. However, instead of pairs of situations, it relates extensional situations s to intensional histories l . Intuitively, $C(l, p, s)$ means that l and s agree on what (physical and sensing) actions were executed, and what the sensing outcomes were. Additionally, it keeps track of the “plausibility offset” p between the intensional and extensional viewpoint. (Note that for any l , there are generally multiple compatible situations s with varying plausibility). The predicate is given by an initial axiom and an SSA:

$$C(l, p, S_0) \equiv l = \epsilon \wedge p = 0 \quad (18)$$

$$C(l, p, do(a, s)) \equiv \exists a', l', p'. l = ido(a', l') \wedge \quad (19)$$

$$\begin{aligned} C(l', p', s) \wedge Apl(a, p'', s) \wedge p = p' + p'' \wedge \\ (SF(a, s) \equiv r(a')) \wedge a = o(a') \end{aligned}$$

We then have a theorem that says that $B(s, n, s')$ is equivalent to $B^i(s'', n', l)$ in case that s and s' are compatible according to Map ; l and s are compatible with offset p according to C ; and n is the sum of n' and p (see Fig. 1):

Theorem 2. Let Σ be a BAT satisfying (11)–(19). Then

$$\Sigma \models B(s, n, s') \equiv \exists s'', p, n', l.$$

$$Map(s, s'') \wedge B^i(s'', n', l) \wedge C(l, p, s') \wedge n = n' + p$$

¹¹We can obtain $o(\cdot)$ and $r(\cdot)$ from (9): for every $ST(a(\vec{x}), s) \equiv \{\neg\}SF(a'(\vec{x}), s)$, assert $a'(\vec{x}) = o(a(\vec{x}))$, and $r(a(\vec{x})) = \top/\perp$ if the negation sign is absent/present.




Figure 1: Relation between B and B^i

Proof. By induction over the length of s' . \square

Finally, our approach captures that of Shapiro et al. (2011) (SPLL for short) as follows. They use a binary B with SSA

$$B(s'', do(a, s)) \equiv \exists s'. \quad (20)$$

$$B(s', s) \wedge s'' = do(a, s') \wedge (SF(a, s') \equiv SF(a, s)).$$

A function $pl(s)$ associates plausibility values to situations s , which remain unchanged by actions:

$$pl(do(a, s)) = pl(s) \quad (21)$$

A formula is then believed if it holds in all B -related situations with minimal plausibility:

$$MP_S(s', s) \doteq B(s', s) \wedge \quad (22)$$

$$\forall s''. B(s'', s) \supset pl(s') \leq pl(s'')$$

$$Bel_S(\phi, s) \doteq \forall s'. MP_S(s', s) \supset \phi[s'] \quad (23)$$

Let Σ_S denote a BAT according to (Shapiro et al. 2011). We can simulate the same behaviour with our framework having our BAT Σ include the (only) initial axiom

$$Init(s') \wedge Init(s) \supset \quad (24)$$

$$[B(s', p, s) \equiv B(s', s) \wedge pl(s') = p]$$

Thus, SPLL's pl value is encoded as the second argument in our 3-place B fluent, and it inherits all initial properties of the 2-place B (such as being transitive and Euclidean). To also ensure compatibility in future situations, we assume

1. preconditions are true everywhere: $Poss(a, s) \equiv \top$;
2. Σ uses the same sensing axioms as Σ_S ;
3. all action plausibilities are zero: $Apl(a, p, s) \equiv p = 0$;
4. no nature's choice arguments (thus $Ieq(a, a') \equiv a = a'$).

Then we have:

Theorem 3. $\Sigma_S \cup \Sigma \models \forall s. Bel_S(\phi, s) \equiv Bel(\phi, s)$.

Note that since SPLL's framework captures (Scherl and Levesque 2003) as a special case, our approach also generalizes the latter.

6 Discussion

The presented framework generalises and integrates work in reasoning about action and belief change in various ways, including an account of nondeterminism which allows physical and sensing actions that may fail or behave in unexpected ways, and an agent model that allows for contingent beliefs

along with beliefs regarding potential alternative states of affairs. Our account of nondeterminism is fundamentally epistemic, in that, while an agent executes a deterministic action, it doesn't have direct access to those action arguments controlled by nature, and so it may be ignorant of the outcomes of the action. This is reminiscent of earlier work on Communicating Sequential Processes (Hoare 1985), where *angelic* nondeterminism can be considered as under the control of an agent (in that a nondeterministic choice is towards termination, if possible), while *demonic* nondeterminism is not under an agent's control, and can be used to model external factors such as sensor readings. Angelic nondeterminism also appears in GOLOG (Levesque et al. 1997), an agent programming language based on the situation calculus, in which an agent may (essentially) nondeterministically select a program to execute or an object as the subject of an action.

With regards to related work, it has been noted that the approach generalises (Shapiro et al. 2011). The approach also developed out of work by Delgrande and Levesque (2012; 2019) (DL) on belief and nondeterministic action, although in the end the approaches differ significantly. First, DL addressed *belief revision* in a dynamic environment. On sensing a fact or being told a proposition, the agent would believe this fact (or proposition) held, while allowing the possibility that it might be false. Plausibilities of situations were modified according to a set revision strategy, in the case of (Delgrande and Levesque 2019) using the specific recipe given in (Darwiche and Pearl 1997). In contrast, here we adopt a "causal" approach, and an axiomatiser must specify exactly what transpires in normal and exceptional cases. Thus for example a light sensor might normally work and gives the correct result, but it may transitorily report *on* or *off*, or give a parity error, each with their own plausibility, and perhaps under specific conditions. (This is why, basically, if there are no exceptional cases we obtain the approach of (Shapiro et al. 2011).) The present approach is simpler than DL; in particular we have a 3-place (as opposed to a 4-place) B fluent and the successor-state axiom for B is significantly simpler. Last, this work significantly generalises DL (and indeed other such approaches in the situation calculus) by giving a complete account of action preconditions and incomplete knowledge, along with a theory of fallible sensing, and the introduction of (agent-centric) intensional belief.

Similar comments apply to (Bacchus et al. 1999) (BHL), which deals with noisy sensors and actuators in the situation calculus. Here, as with other approaches, nontrivial action preconditions require that the agent have complete knowledge of a situation. In BHL the emphasis is narrower than the present approach, addressing *noise* in physical and sensing actions, where an action outcome may deviate probabilistically from the true or expected outcome. In contrast, in the present approach, the emphasis is on action outcome *plausibility*, reflecting nondeterminism and possible action failure. Hence a coin flip will be believed to be either heads or tails, with an implausible outcome being that the flip fails and the coin remains held. While BHL don't address sensor and actuator faults explicitly, it seems feasible to extend their formalism to express probabilistic variants of what we propose in this paper.

We suggest that the framework may shed light on other aspects of reasoning about action, and that it may also be extended in various interesting directions. The approach may be seen as providing a resolution to the *qualification problem* (McCarthy 1977). In our approach, one can specify action preconditions as usual; these are taken as sufficient conditions for an agent to believe that an action can be executed. One can also specify what can go wrong, and if an action happens to fail, this can be seen as some unspecified precondition not being satisfied.

The approach can be straightforwardly modified to encode revision strategies, and it may provide a nuanced means of exploring belief change (Alchourrón *et al.* 1985; Peppas 2008) in a (more) concrete or realistic setting. In the standard semantic account of belief revision, revision by a formula ϕ involves modifying a plausibility ranking so that all minimal worlds in the resulting ranking have ϕ true. The difficulty is that this approach is so abstracted as to provide virtually no guidance as to how revision should actually be carried out. Arguably, any realistic, practical revision operator must take into account real-world phenomena such as where the information came from, how an action may fail, the reliability of sensors, and so on.

Related to this, an obvious item for future work is to address progression and regression. Our approach, as with others in the situation calculus (and indeed most belief revision approaches), is *reflex* (or perhaps *Markovian*) in that, for any action, the resulting total preorder is completely determined by the prior preorder and the action along with its associated plausibility. As pointed out in (Hunter and Delgrande 2011), this is sometimes insufficient for belief change, and in some cases one must regress a formula, do a revision in an earlier situation, and then progress the result.

Our account is qualitative, and there are well-argued reasons for taking this stance (for example, that it is more intuitive and understandable, that it is not clear where probabilities come from, etc.). Of course, there are also well-argued reasons for taking a quantitative, probabilistic approach. It would be of interest to recast the present approach in probabilistic terms; and it seems that there is no great impediment to doing so and that our qualitative approach could be straightforwardly replaced by a probabilistic one. However, a very intriguing related possibility is to *combine* qualitative and probabilistic approaches in the framework, where qualitative plausibility levels are maintained, but within a rank probabilities are employed. Then one could express that, in flipping a coin, normally it will come up heads 50% of the time and tails 50% of the time. Exceptionally, it may land on its side or the action may fail.

In a separate direction we propose to add non-transitory abnormal actions, that is, actuators and sensors behaving predictably but abnormally, that is, they may be malfunctioning. This can be extended to things functioning abnormally in the domain, for example a light being burned out. This would seem to be straightforward to incorporate, using conditional actions, but it would complement the account of transitory action failure. As well, it might be useful for an approach to diagnosis in a dynamic setting.

Last, the approach would seem to provide a good basis for

epistemic planning (Baral *et al.* 2017) in the case of non-determinism, and fallible actuators and sensors. For future work we propose to investigate such a planner in a PDDL dialect, based on the notion of intensional belief, and employing an iterative-deepening-like control structure with respect to plausibilities.

7 Conclusion

We have presented a model of an agent, situated in a (seemingly) nondeterministic environment in which physical and sensing actions may fail or may result in unexpected outcomes. The agent maintains a set of contingent beliefs about the domain that it believes to hold, as well as alternative possibilities that could be the case. Beliefs are modified as sensing and other actions are executed. The approach is developed within an epistemic extension of the situation calculus, incorporating a plausibility ranking over situations to represent the agent's belief state. Two distinct but related notions of belief are developed, an extensional account that models an agent's beliefs wrt actually-executed actions, and an intensional account representing beliefs from the agent's point of view.

Our primary point of departure is that an action can be seen as having two aspects. The intensional aspect roughly corresponds to the standard conception of an action in the situation calculus, and it is that aspect of the action that an agent is aware of and has control over. The extensional aspect is the action that is actually executed in the domain, in the general case beyond direct control of the agent. Consequently the agent's beliefs will evolve according to what it believed it executed, or according to the most normal outcomes of an action, while the world will evolve according to the actions that the agent actually executed. The agent's beliefs can be brought into alignment with reality via sensing, provided that a sensor in fact (usually) works correctly.

This approach extends previous work by providing a complete, qualitative account of reasoning in an unpredictable domain, with potentially unpredictable (or fallible) physical and sensing actions. We retain the results of basic action theories as well as subsuming (Shapiro *et al.* 2011), and so inherit the formal results of these theories. While we present a specific approach, the overall framework is very general; for example it is straightforward to encode a specific approach to belief revision. Consequently the approach may serve as a framework from which to examine general issues in belief change and reasoning about action – for example the interplay of knowledge gained from various action types and differing degrees of reliability. As well, we suggest that it would be of interest to examine epistemic planning from the standpoint of intensional belief in the approach.

Acknowledgements

We thank the reviewers for their helpful and pertinent comments. Hector Levesque originally proposed the notion of action arguments controlled by nature. Financial support was gratefully received from the Natural Sciences and Engineering Research Council of Canada.

References

- C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- 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.
- C. Baral, T. Bolander, H. van Ditmarsch, and S. McIlraith. Epistemic Planning (Dagstuhl Seminar 17231). *Dagstuhl Reports*, 7(6):1–47, 2017.
- R.J. Brachman and H.J. Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers, 2004.
- A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
- G. De Giacomo and H. Levesque. Progression and regression using sensors. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 160–165, Stockholm, Sweden, 1999.
- J.P. Delgrande and H.J. Levesque. Belief revision with sensing and fallible actions. In *Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, Rome, Italy, 2012.
- J.P. Delgrande and H.J. Levesque. An epistemic approach to nondeterminism: Believing in the simplest course of events. *Studia Logica*, 107:859–886, 2019.
- L. Fang and Y. Liu. Multiagent knowledge and belief change in the situation calculus. In *Proceedings of the (AAAI) Conference on Artificial Intelligence*, page 304–312, 2013.
- L. Fang, Y. Liu, and X. Wen. On the progression of knowledge and belief for nondeterministic actions in the situation calculus. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, pages 2955–2963. AAAI Press, 2015.
- P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.
- C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
- A. Hunter and J.P. Delgrande. Iterated belief change due to actions and observations. *Journal of Artificial Intelligence Research*, 40:269–304, 2011.
- T.Q. Klassen, S.A. McIlraith, and H.J. Levesque. Specifying plausibility levels for iterated belief change in the situation calculus. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning*, pages 257–266, 2018.
- T.Q. Klassen, S.A. McIlraith, and H.J. Levesque. Changing beliefs about domain dynamics in the situation calculus. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 572–581, 9 2020.
- G. Lakemeyer and H.J. Levesque. Situations, si! situation terms, no! In D. Dubois, C. Welty, and M.A. Williams, editors, *Proceedings of the Ninth International Conference on the Principles of Knowledge Representation and Reasoning*, pages 516–526, Whistler, BC, 2004. AAAI Press.
- H.J. Levesque, R. Reiter, F. Lin, and R.B. Scherl. Golog: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31:59–83, 1997.
- H.J. Levesque, F. Pirri, and R. Reiter. Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science*, 3(18), 1998.
- H. Levesque. What is planning in the presence of sensing? In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI96)*, pages 1139–1146, 1996.
- J. McCarthy. Epistemological problems in artificial intelligence. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1038–1044, Cambridge, MA, 1977.
- A.C. Nayak, M. Pagnucco, and P. Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146(2):193–228, 2003.
- 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.
- D. Rajaratnam, H. Levesque, M. Pagnucco, and M. Thielischer. Forgetting in action. In *Proceedings of the Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*. AAAI Press, 2014.
- R. Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In V. Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, 1991.
- R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, Cambridge, MA, 2001.
- R.B. Scherl and H.J. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.
- C. Schwing, G. Lakemeyer, and M. Pagnucco. Belief revision and projection in the epistemic situation calculus. *Artificial Intelligence*, 251:62–97, 2017.
- S. Shapiro and M. Pagnucco. Iterated belief change and exogenous actions in the situation calculus. In *Proc. ECAI04*, 2004.
- S. Shapiro, M. Pagnucco, Y. Lespérance, and H.J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.
- 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.