On the Syntax and Semantics of Effect Axioms [1]

Haythem O. ISMAIL
German University in Cairo
Department of Computer Science
haythem.ismail@guc.edu.eg

Abstract. Effect axioms constitute the cornerstone of formal theories of action in AI. They drive standard reasoning tasks, especially prediction. These tasks need not be coupled with actual acting; the reasoning agent is, thus, typically given an *ex post acto* narrative of what actions took place. An *acting* agent, however, has no access to such knowledge; it needs to face what we call the event categorization problem, and figure out what actions it did. Until this is achieved, effect axioms will be useless. A careful review of the literature on effect axioms reveals that their syntax, semantics, and ontological commitments are so deeply entrenched in the armchair reasoning about action paradigm, that they cannot be used in resolving the event categorization problem. By enriching the ontology of action theories, we propose a different approach for representing effects of actions that unifies the two views. The enriched ontology is independently motivated by linguistic concerns.

Keywords. Knowledge representation, actions, events, states.

1. Introduction

Reasoning about action has kept the AI agenda busy for several decades. It has been the origin of such important and nagging problems as the frame problem, the ramification problem, and the qualification problem [1, for instance]. Investigating these problems has culminated in a number of mature action theories that allow reasoning about action to go mostly unhindered [1,2,3, for example].

It is important, however, to distinguish two modes of reasoning about action:

Mode 1. Given knowledge of the executability and effects of some actions, we need to answer questions pertaining to the outcome of executing action instances in a given situation. Note that this is pure armchair reasoning; no *acting* needs to be actually taking place. This reasoning mode is the one with which the theories of action indicated above are mostly concerned.

Mode 2. An agent is acting, and needs to reason about its own actions as it executes them. In this paper, I am not concerned with aspects of such reasoning related

[1] To appear as Haythem O. Ismail, On the Syntax and Semantics of Effect Axioms. In Carola Eschenbach and Michael Grüninger (eds.), *Formal Ontology in Information Systems: Proceedings of the Fifth International Conference (FOIS 2008)*, IOS Press, Amsterdam, 2008. All citations should be to the published version.

to execution monitoring and error recovery in the context of planning. Rather, I focus on the less widely considered aspect related to beliefs of agents about what they are doing, and *have* done, as their own actions unfold in time. Such concerns mainly emerge in the philosophy of action [4,5,6, for example], though AI has occasionally been interested in these questions [7,8].

I will argue that taking the Mode-2 perspective raises problems that are different from those addressed by current Mode-1 theories of action. In particular, armchair reasoning about action naturally assumes that the reasoning agent knows what actions took place (they are given as part of a "narrative"), and the task is to predict outcomes. (Not that it is a simple task.) For an acting agent, however, knowledge of what actions took place is not readily available. Even in the simple, single-agent world usually investigated in AI, the lonely agent needs to reason in order to figure out what it did; it needs to be able to *categorize* its activity before it can figure out the ramifications or even report (possibly in a natural language) on what it did [9]. The reasoning problem of an agent's categorizing its own activities will be referred to as the *event categorization problem*.

Perhaps the central ingredient of theories of action are effect axioms: axioms capturing the causal laws of the domain. A close examination of the literature shows that current treatments of effect axioms are so deeply entrenched in the Mode-1 paradigm that their syntax, semantics, and ontological commitments are not suitable for Mode-2 reasoning. By enriching the ontology of action theories, we propose a slightly different approach to representing the effects of actions that can account both for Mode-1 and Mode-2 reasoning. Incidently, the enriched ontology is independently motivated by linguistic concerns.[2]

Effect axioms are examined in Section 2. Section 3 motivates the event categorization problem. Sections 4 and 5 analyze the inadequacy of effect axioms, in their current form, to account for event categorization. An approach to effect axioms that does this is presented in Section 6.

2. Effect Axioms: Round I

In formal AI action theories, effect axioms are logical formulas which specify, for each action, its *direct* effects. These are effects that cannot be specified by general domain constraints. Table 1 lists some of the common approaches to representing effects of actions. Unbound variables are universally quantified with widest scope. The examples are drawn from the extensively-studied Yale shooting scenario [10].

The main difference between the theories presented in [11], [1], [2], and [3] (other than the logics they employ) is their treatment of major problems, such as the frame problem. However, we shall only concentrate on those differences relevant to the Mode-1/Mode-2 distinction. A generic effect axiom has the following form (see the exemplary axioms of [1] above):

[2]The reasoning task of *explaining* observations is one in which the reasoner does not know all actions that took place. This is related, but not identical, to the issue that concerns us. Explanation is usually achieved by abduction. As we shall see below, we can adhere to mostly-monotonic deduction, without the need for the specialized machinery required by abduction. Alternatively, our approach may provide a way to *focus* abductive reasoning in the special case we are considering.

Reference	**Syntax**	**Examples**
[11]	$C(\mathsf{a}) = \{\eta_i\}$	$C(\mathsf{load}) = \{\mathsf{Loaded}\}$ $C(\mathsf{shoot}) = \{\neg\mathsf{Alive}\}$
[1]	$[Q_i \wedge Q] \supset \mathrm{Initiates}(a, \eta_i, s)$ $[Q_i \wedge Q] \supset$ $\mathrm{Terminates}(a, \eta_i, s)$	$\mathrm{Initiates}(\mathrm{Load}, \mathrm{Loaded}, s)$ $\mathrm{HoldsIn}(\mathrm{Loaded}, s) \supset$ $\mathrm{Terminates}(\mathrm{Shoot}, \mathrm{Alive}, s)$
[2]	$Q_i \supset \eta_i(do(a, s))$ $Q_i \supset \neg\eta_i(do(a, s))$	$\mathrm{Loaded}(do(\mathrm{Load}, s))$ $\mathrm{Loaded}(s) \supset \neg\mathrm{Alive}(do(\mathrm{Shoot}, s))$
[3]	$t + 1 : \eta_i \Leftarrow t : a \wedge Q_i$	$t + 1 : Loaded \Leftarrow t : Load$ $t + 1 : \neg Alive \Leftarrow t : Shoot \wedge Loaded$

Table 1. Examples of effect axioms.

$$[Q_i \wedge Q] \supset R(\eta_i, a)$$

a is an action term, η_i denotes an effect of (the denotation of) a, and $R(\eta_i, a)$ is a formula involving a and η_i. R may explicitly allude to an element of causality ([1] and [3]) or not ([11] and [2]). Shanahan's "Initiates" and "Terminates" indicate what Talmy refers to as *onset-causation* [12]—causation where the cause only *triggers* p; p persists by inertia. Giunchiglia et al's modal $\Leftarrow$ refers to Talmy's *extenet-durational causation* [12]—causation where the persistence of the cause is necessary for the persistence of the effect.[3]

The antecedent of our generic effect axioms, in general, contains two expressions: Q_i and Q. Both qualify the effect axiom but, given the subscript, Q_i somehow depends on η_i and Q depends only on a. The distinction is roughly the following: Q is a conjunction of *executability conditions*, required for a to be executable; Q_i is a conjunction of *effectiveness conditions*, required for a to yield the particular effect η_i.

Interestingly, this distinction does not seem to get the attention it deserves in many theories of action. In [11], for example, *all* qualifications are decoupled from effect axioms, enumerated (akin to the set $C(a)$ of consequences) as part of the action description. As a result, Ginsberg and Smith cannot specify that a loaded gun is needed for the shooting to be lethal, without requiring a loaded gun for the shooting to be possible in the first place.[4] On the other hand, both types of condition appear undistinguished in Shanahan's effect axioms [1]. Thus, (i) having a gun and (ii) the gun's being loaded would appear undistinguished in an effect axiom relating pulling the trigger to death. This makes one wonder about the semantics of action terms; for the absence of either condition will merely block the lethal effect of trigger-pulling. Now, whereas the non-lethal pulling of the trigger of an unloaded gun is a successful pulling of the trigger, no action qualifies as a pulling of the trigger in the absence of a gun.

Reiter [2] (but not Reiter in [13]) and Giunchiglia *et al.* [3] consistently recognize this distinction by providing separate "precondition axioms" specifying executability conditions; effectiveness conditions are conjoined in the antecedent of effect axioms.

[3]This interpretation of $\Leftarrow$ is only revealed by its use in domain constraints in solving the ramification problem.

[4]This might be Hanks and McDermott's original sin, but, for the Yale shooting scenario [10] to make the point it was intended for, everyone assumes that the gun could be successfully shot without being loaded (which is quite revealing as we shall see below). Readers who find this counter-intuitive may replace the act of shooting by that of pulling the trigger.

Now, even if the formalism allows for a distinction between executability and effectiveness (as in [2,3]), it is not clear how this feature will be exploited. In fact, it seems that the distinction between what should go into the effect axioms and what should be separately asserted is, to a big extent, arbitrary. If they do give definitions of qualification, authors typically give vague definitions under which several fine-grained notions are conflated. The following examples come from [13], in the context of database update.

Example 1. It is possible to register a student in a course only if they have passed all pre-requisites.

Example 2. You can only change the grade of a student in a course to grade g if g is different from their current grade.

Example 3. A student may drop a course only if they are registered in it.

It is interesting to note that each of these examples illustrate a different type of qualification. The first is an example of *normative* qualifications; there is nothing physically impossible about registering a student in any course, but one *ought to* only register in courses for which they have all pre-requisites. The last is akin to an executability condition (but see Section 5.2); one cannot remove a student's name from a list if it is not already there. Each of these notions may be further divided into even more fine-grained ones, giving rise to an ontology of qualifications.[5]

The second example is the really interesting one. Why is it necessary for the new grade to be different from the old one? This is not a matter of executability: one can always delete the old grade and then insert it again. Nor is this a matter of normative conventions: there is nothing inappropriate about performing this vacuous update. Nor is it a matter of effectiveness for that matter; for what effect will be blocked if the two grades are identical? The only problem is that we cannot *categorize* the action as a *change* unless the new grade is different from the old one. This being said, we should now turn to event categorization.

3. The Event Categorization Problem

It is always said that there is a striking similarity between the ways we conceive of time and those in which we conceive of space, at least as revealed by language [14,15,16,12, for instance]. In particular, objects (denoted by count terms) correspond to events, and matter (denoted by mass terms) corresponds to states. I will take this fairly acceptable view as a basis for my distinction between what I call "events" and what I call "states".

One can view the conceptual difference between objects and matter as topological: objects are conceived of as topologically-closed and matter as topologically-open. What the distinction amounts to is that objects have their boundaries as parts; matter on the other hand, though always exists in the constitution of bounded objects, is not conceived of as having any boundaries as parts. On this view, any bounded amount of matter constitutes an object. Thus, a pile of sand, a lake, and a beam of light are objects, but sand, water, and light are only matter.

[5] For example, it is probably physically impossible for me to do a double somersault, given my fitness and lack of training. But maybe I can learn. However, I do not think I can ever learn to shoot a gun if I do not have one.

Similarly, events are (temporally) closed situations, ones that have their boundaries as parts, and states are open situations (also see [17]). This loosely corresponds to the linguistic distinction between bounded and unbounded sentences. For example, the imperfective (unbounded) sentence (1a) describes the street-crossing situation as a state, since the temporal boundary of the situation (i.e., its end) is not part of the description; as far as we can tell, the speaker might still be crossing the street. On the other hand, the perfective (bounded) (1b) describes the situation as an event, a bounded whole. Because its temporal boundary is part of an event, event-sentences always imply that the reported situation has come to an end.[6]

(1) (a) I was crossing the street.
    (b) I crossed the street.

Now, as far as language is concerned, the space-time analogy is almost perfect. However, there is a certain respect in which it seems not to hold. In particular, the analogy fails in the way we actually *experience* time and space. In our everyday experience, we encounter objects; we see them, touch them, and (possibly) manipulate them in a, more or less, direct way. However, we rarely encounter matter per se; matter *typically* comes packaged as objects. Thus, we do not see "wood", "glass", or "paper"; we see chairs, bottles, and books.[7]

Our temporal experience, on the other hand, follows the exact opposite pattern. We never experience an event, a *whole* situation; no sooner have we reached the end of a situation, than its beginning has already moved into the past, beyond the reach of our conscious experience. Instead, the world continuously unfolds, presenting us with a continuous flux of states. Evidently, whatever is "now" the case is a state, never an event, for an event has its boundary as an essential part and, thus, can only exist in retrospect, when it has reached an end.[8]

But, if experience consists of only a cascade of states, where do events come from? Events are purely conceptual beasts; we conjure them up by conceptualizing a whole out of some state's starting to hold, holding for a while, and then ceasing. Logically, we must infer event occurrences from patterns of states. One might propose the following: an event has occurred if some state started to hold, held for a while, and then ceased. This is fine and good; the problem is that it only describes the occurrence of some *uncategorized* event token. An uncategorized event token is not very interesting, since one cannot derive any consequences of its occurrence, nor can one report its occurrence in any natural, informative way. The problem then is to infer, not only the occurrence of an event token, but also a categorization thereof. I shall call this *the event categorization problem*.

An acting agent inevitably faces the event categorization problem. It does so in the need to categorize its own acts—the primary example of Mode-2 reasoning.

---

[6] Note that by speaking of boundedness here I am not referring to *telicity*. The two notions are often conflated [18,19, for example], but several authors have distinguished them [20,21, for example]. Boundedness is a purely topological notion; telicity involves goal-directedness or, in general, some notion of a natural boundary. The role of telicity in my proposal will be discussed in detail in Section 5.2.

[7] Or we see *chunks* of wood, glass, and paper; but these are also objects, given their boundaries.

[8] This point has been made by Ismail [8, Ch. 3] and independently by Galton [22].

4. Effect Axioms: Round II

The theories of action presented in Section 2 can happily accommodate agents doing Mode-1 reasoning. However, as they stand, the same theories are not prepared for Mode-2 reasoners. In particular, all effect axioms presented in Table 1 presume that we know what action took place. That is, in order for these axioms—which drive all inferences in an action theory—to be at all useful, the agent has to have categorized its own acts. Otherwise, the agent will not know which effect axioms are applicable.

One might suggest that we can still proceed with existing Mode-1-oriented action theories if we are to give up Mode-2 reasoning all together, and tackle the event categorization problem without resorting to reasoning. In what follows, let $\mathbb{M}(c)$ be the motor program the agent executes whenever it decides to perform an instance of the act category denoted by the term c.[9] We can readily dismiss one obvious loophole:

> *If the agent starts to perform $\mathbb{M}(c)$, then, when it is done, it should believe that its action is of category c.*

The problem, of course, is that $\mathbb{M}(c)$ may fail. For example, in an attempt to shoot a gun, the agent's finger slips in the process of pulling the trigger. While someone might propose that this is indeed a shooting, albeit a failed one, it should be clear that this is just a linguistic trick: a failed shooting is as good a shooting as an alleged murderer is a murderer. Also clearly, if such an act is assumed to be a shooting, then we are committed to deriving all consequences of shooting.

A more realistic suggestion will be the following:

> *Each motor program should end with a sensing step, checking for certain conditions signalling success or failure. Whenever $\mathbb{M}(c)$ succeeds, then an instance of c has been performed. This information may be added to the agent's knowledge base by the motor program itself.*

While this may work for *primitive* act categories, which are directly grounded in motor programs through meta-theoretical association [9], it does not scale up to *high-level* act categories. The latter are associated with motor programs through reductions to primitive acts, expressed in statements of the object language. For example, suppose that pulling the trigger is primitive. It is not hard to construct a corresponding motor program that would sense whether the trigger indeed moved. However, if one advises the agent that it can shoot by pulling the trigger, then either

1. one will risk the agent's having misconceptions about what it did, since a successful pulling of the trigger need not be a successful shooting; or
2. one will have to change the motor program so that it also listens for a bang. (But then a pulling of the trigger would *only* be successful if a bang is heard.)

Both options are clearly unsatisfactory, and any other option will have to involve some sort of reasoning.[10] Further, this suggestion amounts to assuming that the agent can acquire beliefs about event occurrence through proprioception, which goes against our discussion in Section 3.

[9] I am, thus, assuming a first-order theory in which act categories are denoted by terms (similar to "Shoot" and "Load" in Table 1).

[10] Not to mention that sometimes it is not clear what condition to check for within the motor program.

5. Two Important Distinctions

5.1. Necessary and Contingent Effects

Action theories in AI distinguish between direct and indirect effects of actions. Among the direct effects, we propose a distinction between those that are *necessary* and those that are *contingent*. For the agent to categorize an action as c, it has to check for all, and only, necessary effects of c. We can identify necessary effects with effects that are not qualified, and contingent effects with those that are. Qualified effects of actions are only achieved under some conditions. For example, shooting is only lethal if the gun is loaded (see Table 1). Unqualified effects, on the other hand, *necessarily* ensue: loading the gun unconditionally results in its being loaded.

Existing theories of action would, however, have counter-intuitive consequences should we apply the above definition of necessary effects without any provisions. Recall Reiter's Example 1 from Section 2. In Reiter's theory, the student's being registered in a course is not a necessary effect of registering, since the effect axiom is qualified. This is counter-intuitive; it seems that you can only categorize an action as one of *registering* if it results in the student's being *registered*. The trouble with this is probably rooted in the conflation of different types of qualification (see Section 2). In particular, passing all its pre-requisites is a normative qualification for registering in a course. Normative conditions should not be stated as qualifications in effect axioms; they should appear elsewhere in the theory (with special syntax) to indicate when an action is *permissible*.

Similarly, in Reiter's Example 2, it seems that once you've changed a grade to g, then the current grade's being g is an unconditional, necessary effect. The problem is that the old grade's being different is not really a qualification; rather, it has to be different for the performed action to be a *change* of grade.

Thus, we shall define necessary effects of an act category to be unqualified effects, provided that we do not include normative conditions, executability conditions, and the negation of the effect as qualifiers in the effect axiom. In what follows, let $\mathbb{NE}(c)$ denote the set of necessary effects of act category c.

5.2. Telic and Atelic Acts

Consider the following linguistic reports.

(2) (a) I ran.
    (b) I ran to the store.
    (c) I ran toward the store.
    (d) I ran past the store.

For our acting agent to be able to honestly report (2a), it only needs to have been running for a while. Thus, if the state of running ever held, an event of running did occur. This property is characteristic of event categories commonly referred to in the literature as *atelic* [15, for instance]. For the *telic* category reported by (2b), (i) the agent has to have run, (ii) the running has to have stopped at the store, and (iii) the running has to have caused the agent to be at the store. Thus, in order for the agent to form a belief along the lines of (2b), it needs to monitor what it is doing, making sure that it has culminated in a certain way, and to involve some sort of causal reasoning. Clearly, a final state in which a telic event naturally culminates is a necessary effect.

Beyond the standard telic/atelic distinction, [8] points out event categories with different occurrence conditions. Sentence (2c) reports an occurrence of a *left-atelic* event category: similar to (2b), there *is* a natural final state (being at the store), but it should not be reached. (2d) is a report of a *right-atelic* event category: a necessary effect has to be achieved, but the agent's activity may continue beyond that achievement.[11] For limitations of space, I cannot embark on a detailed analysis of this ontology of event categories. A formal treatment may be found in [8]. I will henceforth stick to the tradition, and consider only telic and atelic categories.

Now, let us take some time to discuss what we mean by telicity.[12] The above discussion, like traditional discussions of telicity [23,24, for example], gives the impression that telicity is merely about a natural *end* point. However, several authors [15, for example] have pointed out that telicity is not just a matter of states that hold in the aftermath of an event. This is particularly true for examples like Example 2 of Section 2, where changing a grade (a telic event) can only occur if it effects a transition from an initial, pre-action state to a final, post-action state.

A detailed analysis of telicity may be found in [15]. However, for the sake of the relatively modest objective of this paper, identifying occurrence conditions of telic event categories need not require such an ambitious analysis. In particular, I will take telic events to be effecting a transition from an initial state s_i, that holds at the onset of the event, to a terminal state s_t that is caused by it. In typical examples of telicity, s_i and s_t are contradictory states. For example, in (2b), s_i is the state of not being at the store and s_t is being at the store. Thus, not only should s_t hold at the end of the event, but it should also *not* hold immediately before and, crucially, throughout the event occurrence.

In other cases, however, s_i and s_t need only be contrary states. In this case, s_i implies, but is not necessarily identical to, the negation of s_t. For example, consider (3), where it is not sufficient for me end up at the store. In addition, I have to have started running at the park, which is contrary, but not contradictory, to my initially being at the store.

(3) I ran from the park to the store

In general, however, s_i and s_t could be any two states. In fact, the two states may be identical, as in the example of running around a track. But note that in all cases, through out the event occurrence, s_t does not hold. If it does, then we have a case, not of telicity, but of right-atelicity. Whenever s_i is not indicated by the event description, it seems that it is always taken to be the negation of s_t—the weakest contrary-to-s_t state.

Thus, similar to the set $\mathbb{NE}(c)$ of necessary effects of event category c, there is a set $\mathbb{NI}(c)$ of *necessary initial conditions*, containing all the s_is. In most cases, the conditions in $\mathbb{NI}(c)$ are the negations of those in $\mathbb{NE}(c)$. Conditions in $\mathbb{NI}(c)$ are somewhat similar to executability conditions in that they (i) are implied by the occurrence of an instance of c and (ii) need to be achieved by the agent before attempting an execution of c. The

[11] We can further categorize telic acts into those with immediate effects and those with delayed effects (see [8]).

[12] This discussion was motivated by the criticisms levelled by an anonymous reviewer of FOIS-08 at a draft of this paper.

two types of condition are different in that necessary effects *cannot* be achieved if executability conditionas are not satisfied; they could be achieved, however, even if initial conditions do not hold (witness (3)). One could say that executability conditions are conditions on motor programs, whereas necessary initial conditions are conditions on event descriptions.

6. A Unified Framework

In this section, a framework for mode-2 reasoning is presented. I should stress that what will be presented is just a *framework*, not a complete axiomatic system. A complete system needs more space for presentation, and more time for working out all the details. The purpose of the framework is to provide a starting point, and to illustrate the rich epistemological ontology needed for a mode-2 reasoner (as opposed to the relatively coarse ontology of mode-1 reasoning systems).

6.1. Formal Machinery

We shall need a formal language to talk about states, events, and event categories. Only informal semantics will be provided, however. In my so doing, I trust that the reader will not end up in confusion and that nothing is at stake, regarding the point I am trying to make. In addition to a formal semantics, axioms ruling out unintended models will also be needed should this language be actually used for reasoning.

The purpose of the language should also be clear. Statements of the language represent beliefs of a reasoning and acting agent; they do not represent a third-person perspective of what the agent is doing, was doing, or has done. Thus, no terms denoting agents necessarily appear in the language; only a constant I is necessary, denoting the reasoning agent's self-concept [9]. Moreover, statements representing beliefs about what the agent *is doing* are not derived as a result of, say, plan recognition. Beliefs about what it is doing come form the agent's first-person access to its own intentions and bodily feedback (see [5,6] and particularly [4, p. 23]).

The language is a first-order language akin to that of [25], which is a revised version of Allen's interval-based theory [7], where instants are independently included in the ontology. The ontology comprises ordinary objects and individuals, time instants, time intervals, states, event categories, and event tokens. The following predicate symbols are part of our language. (Superscripts indicate adicity.)

- $\sqsubset^2$: $[\![t_1 \sqsubset t_2]\!]$ is true whenever interval $[\![t_1]\!]$ is a proper subinterval of interval $[\![t_2]\!]$
- $\prec^2$: $[\![t_1 \prec t_2]\!]$ is true whenever time $[\![t_1]\!]$ wholly precedes time $[\![t_2]\!]$
- Begins^2: $[\![\mathsf{Begins}(i,t)]\!]$ is true whenever instant $[\![i]\!]$ limits interval $[\![t]\!]$ at its beginning.
- Ends^2: $[\![\mathsf{Ends}(i,t)]\!]$ is true whenever instant $[\![i]\!]$ limits interval $[\![t]\!]$ at its end.
- Within^2: $[\![\mathsf{Ends}(i,t)]\!]$ is true whenever instant $[\![i]\!]$ is within interval $[\![t]\!]$.[13]
- $\mathsf{HoldsAt}^2$: $[\![\mathsf{HoldsAt}(s,i)]\!]$ is true whenever state $[\![s]\!]$ holds at instant $[\![i]\!]$.
- Occurs^2: $[\![\mathsf{Occurs}(e,t)]\!]$ is true whenever event token $[\![e]\!]$ occurs on time $[\![t]\!]$.[14]

[13] An axiomatization of Within may be found in [25].

[14] For simplicity, I am overloading Occurs to work for both durative and instantaneous events; [25] uses two different predicate symbols.

- Cat^2: $[\![\mathsf{Cat}(e,c)]\!]$ is true whenever event token $[\![e]\!]$ is an instance of event category $[\![c]\!]$
- Caused^2: $[\![\mathsf{Caused}(e_1,e_2)]\!]$ is true whenever event token $[\![e_1]\!]$ caused event token $[\![e_2]\!]$.

For the sake of brevity, we define the following convenient predicates:

- $t_1 \supset\!\subset t_2 =_{\mathrm{def}} t_1 \prec t_2 \wedge \neg\exists t_3[t_1 \prec t_3 \wedge t_3 \prec t_2]$[15]
- $\mathsf{Holds}(s,t) =_{\mathrm{def}} \forall i[\mathsf{Within}(i,t) \supset \mathsf{HoldsAt}(s,i)]$
- $\mathsf{MHolds}(s,t_1) =_{\mathrm{def}} \mathsf{Holds}(s,t_1) \wedge \neg\exists t_2[\mathsf{Holds}(s,t_2) \wedge t_1 \sqsubset t_2]$

In addition, we have five function symbols:

- $\uparrow^1$: $[\![\uparrow s]\!]$ is the event category of onsets of state $[\![s]\!]$.
- $\downarrow^1$: $[\![s \downarrow]\!]$ is the event category of cessations of state $[\![s]\!]$.
- $\dot{\neg}^1$: $[\![\dot{\neg}\, s]\!]$ is the unique state that holds at every instant at which $[\![s]\!]$ does not hold.
- Prog^1: $[\![\mathsf{Prog}(c)]\!]$ is the unique state that holds whenever event category $[\![c]\!]$ is in progress
- Clos^2: $[\![\mathsf{Clos}(s,t)]\!]$ is the event token of state $[\![s]\!]$ maximally holding throughout time $[\![t]\!]$

Clos is closely related to MHolds:

AMC. $\mathsf{Occurs}(\mathsf{Clos}(s,t),t') \equiv (t' = t) \wedge \mathsf{MHolds}(s,t)$

Thus, every time a state maximally holds, the agent may easily infer the occurrence of some event, namely the closure of that state at the said time. As pointed out in Section 3, such events are not very useful because they do not fall under any natural category.[16]

The semantics of the progressive operator is, no doubt, mysterious. However, I will content myself with the informal gloss given above, and refer the reader to [27]. Nevertheless, note that I am mainly concerned with acting agents, for which progressive states are primarily experienced while acting. Knowledge of such states is a first-person privilege that is investigated in depth in the philosophy of action [4,5,6]. For how such knowledge may be induced in the case of a robot, see [9]. Essentially, $\mathsf{Prog}(c)$ holds whenever the agent is executing $\mathbb{M}(c)$, for primitive acts, or a plan for performing c, for high-level acts. I will, hence, take it as unproblematic that agents can form beliefs such as $\mathsf{Holds}(\mathsf{Prog}(\mathsf{Run}(\mathsf{I})),t)$ or $\mathsf{Holds}(\mathsf{Prog}(\mathsf{RunTo}(\mathsf{I},\mathsf{Store})),t)$. I shall also restrict Prog by the following axiom.

AP. $[\mathsf{Occurs}(e,t) \wedge \mathsf{Cat}(e,c)] \supset \mathsf{MHolds}(\mathsf{Prog}(c),t)$

This axiom licences inference from statements like (1b) to statements like (1a). Thus, I am not assuming that a notion of intention is an essential ingredient in the retrospective use of Prog.

It should be pointed out, however, that **AP** is not appropriate for all classes of event categories. In particular, it does not hold for categories that are *indefinitely-specified* [28]. Indefinitely-specified event categories are ones that involve an indefinite entity. For

[15] The reader will note that $\supset\!\subset$ is Allen's [7] `meets`. The symbol "$\supset\!\subset$" is the one used for the same relation in discourse representation theory [26].

[16] If we adopt Galton's Po operator [27], then $\mathsf{Clos}(s,t)$ is of category $Po(s)$.

example, the category of (bomb) explosions seems to involve an indefinite bomb, the category of concerts involves an indefinite performer, and the category of my picking up *a* block involves an indefinite block. Indefinitely-specified event categories do not conform to **AP** since multiple tokens of them may overlap in time. Consider the category pick-up-a-block of picking up an indefinite block.[17] Suppose there are two blocks: A and B. I start picking up A with my right hand at time instant i_1. While still in the process of picking up A, I start picking up B with my left hand at time instant i_2. Picking up A ends at i_3, and picking up B ends, later, at i_4. Clearly, the categories of picking up A and picking up B satisfy **AP**. However, pick-up-a-block does not. The reason is that, whereas a token of pick-up-a-block does occur over the interval extending from i_1 to i_3, Prog(pick-up-a-block) does not *maximally* hold over this interval. Rather, it maximally holds over the interval extending from i_1 to i_4.

The above notwithstanding, for the purpose of this paper, I will consider only definitely-specified event categories. In banishing indefinitely-specified categories, I am adopting the common tacit (and possibly sub-conscious) policy of most authors. Nevertheless, a careful examination of indefinitely-specified event categories is doubtlessly called for. This is particularly pressing since this class of event categories seems to challenge many of our intuitions about events.[18]

6.2. Occurrence Conditions

The occurrence conditions of an event category will be represented by statements of the following form, where ϕ is the condition.

- $\exists e[\mathsf{Occurs}(e,t) \wedge \mathsf{Cat}(e,c)] \equiv \phi$

The exact form of ϕ depends on the type of category we are considering. First, events are either instantaneous or durative. I will take instantaneous events to be onsets or cessations of states. To infer the occurrence of an onset, the agent should experience a state not holding, followed by an experience of the same state holding. Similarly for cessations.[19] [20]

AO. $\exists e[\mathsf{Occurs}(e,i) \wedge \mathsf{Cat}(e,\uparrow s)] \equiv$
$$\exists t_1, t_2[\mathsf{Holds}(\dot{\neg}\, s, t_1) \wedge \mathsf{Holds}(s, t_2) \wedge \mathsf{Ends}(i, t_1) \wedge \mathsf{Begins}(i, t_2)]$$

AC. $\exists e[\mathsf{Occurs}(e,i) \wedge \mathsf{Cat}(e, s\downarrow)] \equiv$
$$\exists t_1, t_2[\mathsf{Holds}(s, t_1) \wedge \mathsf{Holds}(\dot{\neg}\, s, t_2) \wedge \mathsf{Ends}(i, t_1) \wedge \mathsf{Begins}(i, t_2)]$$

Durative events are those described in Section 3 as comprising a state starting to hold, holding for a while, and then ceasing. Inferring the occurrence of such events is not

---

[17] See [28] for a more elaborate characterization of the syntactic structure of terms denoting indefinitely-specified event categories.

[18] See [28] for examples of such intuitions.

[19] Events such as winking or hiccupping, which are usually considered punctual ([24]), are not instantaneous on my view.

[20] Since states need to be *experienced* in order for an agent to infer the occurrence of their onsets and cessations, such states cannot hold at isolated instants. This assumption is made explicit in axioms **A0** and **A1** below. In the terminology of [25], these states are "states of motion". States that can hold at isolated instants (Galton's "states of position") cannot be experienced; similar to events, their holding at an instant has to be inferred.

as simple as in the case of instantaneous events; it depends on whether the event is telic or atelic.

For atelic categories, we have the following general condition:[21]

AA. $\mathsf{Cat}(\mathsf{Clos}(\mathsf{Prog}(c), t), c)$

Thus, the closure of the progressive state of an atelic event category is always an instance of it. Using **(AMC)**, (**AP**), and **AA**, it is easy to show that

TA. $\exists e[\mathsf{Occurs}(e, t) \wedge \mathsf{Cat}(e, c)] \equiv \mathsf{MHolds}(\mathsf{Prog}(c), t)$

For telic events, the situation is more complex:

AT1.
$$[\mathsf{MHolds}(\mathsf{Prog}(c), t_1) \wedge \bigwedge_{s \in \mathbb{NE}(c)} \mathsf{Holds}(\dot{\neg}\, s, t_1) \wedge \bigwedge_{s \in \mathbb{NI}(c)} \exists t_2[t_2 \supset\!\subset t_1 \wedge \mathsf{Holds}(s, t_2)] \wedge \bigwedge_{s \in \mathbb{NE}(c)} \exists t_3[t_1 \supset\!\subset t_3 \wedge \mathsf{Holds}(s, t_3) \wedge \mathsf{Caused}(\mathsf{Clos}(\mathsf{Prog}(c), t_1), \mathsf{Clos}(s, t_3))]] \supset \mathsf{Cat}(\mathsf{Clos}(\mathsf{Prog}(c), t_1), c)$$

That is, whatever the agent did on t_1 is an instance of c if, as far as it knows, the necessary initial conditions held as it started c-ing, it was c-ing throughout t_1, the necessary effects of c did not hold throughout t_1 but started as the agent's activity halted, and the performance of that activity is what caused the necessary effects. I believe that the biggest challenge in all this is the final causal link; ultimately, a full analysis of causation along the line of [29,30, for example] is needed. However, this might be a situation in which defeasibly inferring this causal link would be appropriate.[22]

We can now state the following stronger condition on telic event occurrence.

AT2.
$$\exists e[\mathsf{Occurs}(e, t_1) \wedge \mathsf{Cat}(e, c)] \equiv [\mathsf{MHolds}(\mathsf{Prog}(c), t_1) \wedge \bigwedge_{s \in \mathbb{NE}(c)} \mathsf{Holds}(\dot{\neg}\, s, t_1) \wedge \bigwedge_{s \in \mathbb{NI}(c)} \exists t_2[t_2 \supset\!\subset t_1 \wedge \mathsf{Holds}(s, t_2)] \wedge \bigwedge_{s \in \mathbb{NE}(c)} \exists t_3[t_1 \supset\!\subset t_3 \wedge \mathsf{Holds}(s, t_3) \wedge \mathsf{Caused}(\mathsf{Clos}(\mathsf{Prog}(c), t_1), \mathsf{Clos}(s, t_3))]]$$

The right-to-left direction follows from (**AT1**) and (**AMC**). The left-to-right direction could be replaced by a set of conditionals. The first, with $\mathsf{MHolds}(\mathsf{Prog}(c), t_1)$ as a consequent, is just (**AP**). The rest correspond to unqualified effect axioms and axioms for necessary initial conditions.[23]

[21] Once again, this condition does not apply to indefinitely-specified event categories.

[22] On another note, we can assert the causal relation, not between the closures of states, but from the closure of the progressive state to the *onset* of necessary effects. This would be closer to the spirit of Shanahan's "Initiates" [1]. This move will actually be required if we are to explicitly dismiss unbounded intervals from our ontology. For, in that case, necessary effects could not possibly be permanent.

[23] If these axioms are part of our theory, (**AT2**) would be a theorem.

The occurrence condition for the act of changing student a's grade to g (see Example 2 in Section 2) may be stated as follows (assuming pretend-it-is-English semantics).

$$\exists e[\mathsf{Occurs}(e,t_1) \wedge \mathsf{Cat}(e, \mathsf{Change}(\mathsf{grade}(a), g))] \equiv$$
$$[\mathsf{MHolds}(\mathsf{Prog}(\mathsf{Change}(\mathsf{grade}(a), g)), t_1) \;\wedge\; \mathsf{Holds}(\dot{\neg}\, (\mathsf{grade}(a) = g), t_1) \wedge$$
$$\exists t_2[t_2 \supset\!\subset t_1 \;\wedge\; \mathsf{Holds}(\dot{\neg}\, (\mathsf{grade}(a) = g), t_2)] \wedge$$
$$\exists t_3[t_1 \supset\!\subset t_3 \;\wedge\; \mathsf{Holds}(\mathsf{grade}(a) = g, t_1), t_3) \wedge$$
$$\mathsf{Caused}(\mathsf{Clos}(\mathsf{Prog}(c), t_1), \mathsf{Clos}(\mathsf{grade}(a) = g, t_1), t_3))]]$$

Using the above occurrence conditions, an agent can categorize its own acts. It only needs to experience the necessary effects, but all contingent effects and ramifications may be inferred through the standard axioms. Note that whether **(AT2)** or **(TA)** is applicable depends on the teleological features of c. Knowledge of telicity may be explicitly given to the agent, by asserting of durative event categories whether they are telic or atelic. However, it may also be derived through a fine-grained analysis of telicity in the spirit of [15].

7. Conclusions

By distinguishing necessary and contingent effects, and telic and atelic act categories, we presented a language which seems to be expressive enough to serve as a tool for both Mode-1 and Mode-2 reasoning. The underlying ontology makes distinctions that have long been recognized as necessary for natural language semantics [27,24,15].

What has been presented is by no means a theory, it is a framework within which more investigations should follow. The following are some of the issues of which this paper could only scratch the surface.

1. What is a precondition? We distinguished executabiltiy conditions, normative qualifications, and effectiveness conditions. However, we also noted a strong relation between executability conditions and the necessary initial conditions appearing in general occurrence axioms. Where do these fit? Are there other types of preconditions?
2. How rich is the notion of telicity? We distinguished telic and atelic events, but we noted the existence of other classes (left-atelic and right-atelic, telic with immediate or delayed effects, etc.). A careful examination of telicity as a fundamental notion for event occurrence is called for.
3. What is the effect, on the presented system, of admitting indefinitely-specified events? As noted in the paper, at least two of our fundamental axioms will no longer be valid.

To conclude, here are some remarks on constructing formal theories that accommodate both Mode-1 and Mode-2 reasoners:

1. Executability conditions and normative qualifications should be separated from effect axioms.
2. The negation of an effect should not appear as a qualification on the same effect. Rather, it is a *necessary* initial condition, akin to executability conditions.

3. Necessary effects of telic categories should appear only in occurrence conditions; the corresponding effect axioms logically follow.

Alternatively, given a standard Mode-1 action theory that observes the first two remarks, we can always generate occurrence conditions by realizing that telic acts are those with unqualified effects. Such effects are the necessary effects needed for occurrence conditions.

References

[1] Murray Shanahan. *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. MIT Press, Cambridge, MA, 1997.

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

[3] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 153:49–104, 2004.

[4] Kevin Falvey. Knowledge in intention. *Philosophical Studies*, 99:21–44, 2000.

[5] George Wilson. Proximal practical foresight. *Philosophical Studies*, 99:3–19, 2000.

[6] David Velleman. *The Possibility of Practical Reason*. Oxford University Press, Oxford, 2000.

[7] James Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154, 1984.

[8] Haythem O. Ismail. *Reasoning and Acting in Time*. PhD thesis, University at Buffalo, The State University of New York, 2001.

[9] Stuart C. Shapiro and Haythem O. Ismail. Anchoring in a grounded layered architecture with integrated reasoning. *Robotics and Autonomous Systems*, 43(2-3):97–108, 2003.

[10] Steve Hanks and Drew McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33:379–412, 1987.

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

[12] Leonard Talmy. *Toward a Cognitive Semantics*. MIT Press, 2000.

[13] Raymond Reiter. Proving properties of states in the situation calculus. *Artificial Intelligence*, 64(2):337–351, 1993.

[14] Emmon Bach. The algebra of events. *Linguistics and Philosophy*, 9:5–16, 1986.

[15] Manfred Krifka. The origins of telicity. In Susan Rothstein, editor, *Events and Grammar*, pages 197–235. Kluwer Academic Publishers, Dordrecht, 1998.

[16] Ray Jackendoff. Parts and boundaries. *Cognition*, 41:9–45, 1991.

[17] Fabio Pianesi and Achille Varzi. Events, topology, and temporal relations. *The Monist*, 79(1):89–116, 1996.

[18] Renaat Declerck. Aspect and the bounded/unbounded (telic/atelic) distinction. *Linguistics*, 17:761–794, 1979.

[19] Östen Dahl. On the definition of the telic-atelic (bounded-nonbounded) distinction. In Philip Tedeschi and Annie Zaenen, editors, *Tense and Aspect*, volume 14 of *Syntax and Semantics*, pages 79–90. Academic Press, New York, 1981.

[20] Ilse Depraetere. On the necessity of distinguishing between (un)boundedness and (a)telicity. *Linguistics and Philosophy*, 18:1–19, 1995.

[21] Jean-Pierre Koenig and Nuttanart Muansuwan. How to end without ever finishing: Thai semi-perfective markers. *Journal of Semantics*, 17:147–184, 2000.

[22] Antony Galton. On what goes on: The ontology of processes and events. In Brandon Bennett and Christiane Fellbaum, editors, *Formal Ontology in Information Systems: Proceedings of the Fourth International Conference (FOIS 2006)*, pages 4–11, Baltimore, 2006. IOS Press.

[23] Zeno Vendler. Verbs and times. *The Philosophical Review*, 66(2):143–160, 1957.

[24] M. Moens and M. Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14:15–28, 1988.

[25] Antony Galton. A critical examination of Allen's theory of action and time. *Artificial Intelligence*, 42:159–188, 1990.

[26] Hans Kamp and Uwe Reyle. *From Discourse to Logic*. Kluwer Academic, Dordrecht, 1993.

[27] Antony Galton. *The Logic of Aspect*. Clarendon Press, Oxford, 1984.

[28] Haythem O. Ismail. Simultaneous events and the “once-only” effect. In Brandon Bennett and Christiane Fellbaum, editors, *Formal Ontology in Information Systems: Proceedings of the Fourth International Conference (FOIS 2006)*, pages 143–154, Baltimore, 2006. IOS Press.

[29] David Lewis. Causation. *The Journal of Philosophy*, 70(17):556–567, 1973.

[30] Charles Ortiz. Explanatory update theory: Applications of counterfactual reasoning to causation. *Artificial Intelligence*, 108(1–2):125–178, 1999.