Possible-World and Multiple-Context Semantics for Common-Sense Action Planning

Maria J. Santofimia,
Computer Architecture and Network Group,
School of Computing Science.
University of Castilla-La Mancha, Spain

Scott E. Fahlman,
Language Technologies Institute,
Carnegie Mellon University,
Pittsburgh, PA 15213, USA

Francisco Moya, and **Juan Carlos Lopez**
Computer Architecture and Network Group,
School of Computing Science.
University of Castilla-La Mancha, Spain

Abstract

Event management and response generation are two essential aspects of systems for Ambient Intelligence. This work proposes handling these issues by means of an approach with which to model and reason about actions and events which, under the umbrella of a philosophical and common-sense point of view, describes what actions and events are, how they are connected, and how computational systems should consider their meaning. This work uses the Scone Knowledge-Base (KB) system with which to both reason about and model the context and the related events. This approach is capable of generating ad-hoc responses, in terms of actions to be performed, supported by the knowledge about the possible-world and multiple-context semantics.

1 Introduction

It is a well known fact that intelligent systems struggle with innovation and change whereas humans seem to perform well in most cases. Why is this? or what lies beneath this human skill? The response of cognitive science to these questions points out the human ability to handle and reason about *possible worlds*. The notion of possible worlds is used here to refer to those states of affairs or "worlds" which, given an event or a premise, are true in all the worlds considered possible. For example, to state an analogy with the Sherlock Holmes stories, the true facts are provided by the clues in the case. Holmes therefore considers all the *worlds* in which the given premises are true. Note how new clues might lead Holmes to reject worlds that were previously considered to be plausible.

Closely related to the notion of possible worlds, the *context* concept is here understood as the set of facts or propositional knowledge that describe a specific state of the world, in the same way that J. Allen's refers to the *world* concept in [Allen, 1984]. This concept is represented by a description set of both the static and dynamic aspects of the world, thus modeling what is known about the past, present, and future. The J. Allen nomenclature can be used to state that the static aspects of the world are easily captured as *properties* while the dynamic aspects are captured as *occurrences* or *events*.

The notion of *multiple contexts* is connected with that of possible worlds and refers to the mechanism used to concurrently handle the possible-world semantics, at the knowledge-base level. The multiple-context mechanism provides a mean to model actions and events by describing the state of the world before, during, and after the action or event takes place. For example, a `person moving` event gives rise to a new world-state in which the person that moves changes location. If a person moves from the kitchen to the living room, the world-state, before the event takes place, is described by the person being present in the kitchen, while the world-state after the event has taken place is described by the fact that the person is then located in the living room. However, if that person, before moving, approaches an object and picks it up, where is the object after the moving event? Moreover, what will happen if it is a slippery object? The purpose of this work is to model and reason about actions and events, while considering those scenarios that involve the inference of implicit, non-deterministic or delayed effects of events. The following scenarios, extracted from [Mueller, 2006], illustrate those situations that require special attention:

1. In the kitchen, Lisa picked up the newspaper and walked into the living room.
2. Lisa put a book on a coffee table and left the living room. When she returned, the book was gone.
3. Jamie walks to the kitchen sink, puts the stopper in the drain, turns on the faucet, and leaves the kitchen.
4. Kimberly turns the fan's power switch to "on".

In the first scenario, it is easily inferred that since Lisa was intially in the kitchen, she picked up the newspaper while she was there and then took it into the living room. It is also obvious to us that if Lisa is in the kitchen she cannot be in any other room at the same time, since we are considering rooms as non-overlapping spaces in a house. With regard to the second scenario, we can easily infer that if Lisa left the living room, she is no longer there, and that if the book is not there when she returns, something must have happened because things tend to remain in the state they are unless a partiuclar event affects them. The "*frame problem*" concerns determining those things that can be assumed to stay the same from one moment to another. In the third scenario we easily conclude that, after a while, the water will start spilling onto the floor. Finally, with regard to the question of what will happen in the fourth scenario, we can assume that if everything works as it is supposed to, the fan will start up.

1.1 Action Planning in Ambient Intelligence

The objective of this work is to propose an approach for action planning with endowed capabilities to handle the non-trivial aspects of common-sense reasoning. The innovative aspect of this work lies in the heuristics provided by common-sense knowledge concerning actions and events captured in the proposed model.

This work focuses its attention on planning in Ambient Intelligence. Note that Ambient Intelligence environments are characterized by: a) the multiple sources of change affecting the context; b) the device availability aspects that cannot be determined beforehand; and c) the expection of intelligent and autonomous reactions in response to context changes. These aspects, along with the nonlinearity of the problems involved in Ambient Intelligence, are responsible for the small amount literature found in the field.

The strategy followed here consists of: a) proposing a model for actions and events that captures the common-sense knowledge involved; b) representing possible worlds by means of a context activation scheme; c) modeling actions and events in terms of the multiple contexts that describe the world before, during, and after the action or event takes place; d) and finally, rather than considering primitive and compound tasks, in an HTN-like style (Hierarchical Task Network) [Erol *et al.*, 1994], we consider actions that are provided by services and those which are not. By doing this, the proposed approach addresses the device dynamism that characterizes Ambient Intelligence environments.

The remainder of this paper is organized as follows: First, in Section 2 a model for actions and events is proposed and formalized. Section 3 describes how the proposed model is represented in Scone, emphasizing the multiple-context and context activation scheme. Section 4 demonstrates how the key issues of common-sense have been addressed. Section 5 presents an action planning strategy with common sense. A proof of the benefits derived from considering common-sense knowledge as a constituent part of an action planning approach is demonstrated with a case scenario. Finally, Section 6 shows the conclusions drawn from the work presented herein.

2 Modeling actions and events

Actions and events have commonly been treated as being equivalent, or as having the slight difference of considering actions as events which have been intentionally generated [Hommel *et al.*, 2001]. On the contrary, the theory of action for multi-agent planning [Georgeff, 1988] advocates for a distinction between actions and events, although it hints that actions are accomplished by agents in their endeavor to achieve a goal.

Davidson's theories, particularly those regarding the philosophy of action, also identify actions with events, as is argued in [Davidson, 1963]. Actions are described as a combination of two views. On the one hand, actions can be seen as causal explanations of body movements and on the other hand, actions can also be seen as the justifying reason that leads the action to take place. Davidson considers events to be equivalent to actions. The sole difference is that when an action is considered as an event, it is re-described in terms of its effects.

The model proposed here for actions and events adopts the Davidsonian view. It should be highlighted that Cyc [Lenat, 1995], through its language CycL, represents actions and events using a Davidsonian approach. Actions are described as events but are carried out by an agent. The approach implemented in Scone has been extended to include the notion of primary reasons for an action, along with its temporal and location aspects.

Apart from the concept of action and event that concern us here, some other relevant entities must also be considered in relation to actions and events so as to capture their semantics. The following definitions state the foundation of the proposed model for actions and events:

Definition 1. A Context is a set C composed of statements which, when used together, describe knowledge about the world. There may be multiple contexts describing each of the different views of the world. The meaning or truth value of a statement is a function of the context in which it is being considered.

The function $meaning : T, C \rightarrow M$, where T is the set of statements describing the world, C is the set of possible contexts, and M the set of possible meanings, $meaning(t, c)$ therefore returns the meaning or truth value of the statement t in the context c. This can be formally stated as:

$$\begin{aligned} \forall c_i \in C \forall t_i \in T : \\ m_i = meaning(t_i, c_i) \iff t_i \subseteq c_i \end{aligned} \quad (1)$$

The meaning or truth value of a given statement depends on the contexts in which it has been declared.

Definition 2. An Action A is causally explained from the perspective of their relation to the primary reason that rationalizes them. The function $AG : A \rightarrow G$, such that A is the actions, G is the agent, and the function AG returns the agent performing the given action. Furthermore, the function $PR : A, G \rightarrow E$ is the primary reason for an agent performing an action ton seek the effects of the event caused. Finally, the function $PA : A, O \rightarrow G$, such that O is the object, and the function returns the agent that performs the action upon the given object.

$$\begin{aligned} \exists g \in G \exists a \in A \exists o \in O : \\ (AG(a) \wedge PR(a, g)) \iff PA(a, o) \end{aligned} \quad (2)$$

Therefore, an action is performed upon an object, if and only if there exists an agent with a primary reason to perform the action.

Definition 3. An Event E is the individual occurrence that causes changes in the world. The criteria followed by the Davidsonian doctrine on individuation of events argues for the equality of events when the same effects occur. The Davidsonian view is here adapted to internalize the multiple contexts approach. In this paper it is therefore considered that two events are equivalent when the same effects are caused by different actions. The effects of events are captured in the *after context*, while the preconditions for an event to take place are described by the *before context*. The functions $BC : E \rightarrow C$ and $AC : E \rightarrow C$, such that $BC(e)$ and $AC(e)$ respectively return the statements of which the before and after context of a given event are composed. Furthermore, the function $effect : A, O \rightarrow S$, such that S represents the set of statements that describe the world after the event took place.

$$\forall e \in E : (BC(e) \cup effect(a, o)) \rightarrow AC(e) \quad (3)$$

Given the events e_1 and e_2, it can be said that e_1 is equivalent to e_2 when e_2 originates, at least the same effects that characterizes the *after context* of the e_1:

$$\exists e_1, e_2 \in E : e_1 = e_2 \iff AC(e_1) \subseteq AC(e_2) \quad (4)$$

Definition 3. A Service S is provided by a device D and it performs a set of actions upon an object or a set of objects. The function $PD : S \rightarrow D$, such that D is the set of available devices, and the function returns the device or devices that provide a given service.

$$\begin{gathered} \exists s \in S \exists d \in D \exists a \in A \exists o \in O : \\ (PA(a, o) \wedge PD(s)) \rightarrow AG(a) = d \end{gathered} \quad (5)$$

The definition of service therefore implies that the agent of an action provided by a service is a device.

Definition 4. An Object is the set O of possible environmental objects upon which actions are performed. The function $OA : A \rightarrow O$ returns the set of possible objects that can receive a given action.

$$\exists o \in O \exists a \in A \exists e \in E : OA(a) \wedge PA(a, o) \rightarrow e \quad (6)$$

The occurrence of an event e implies the existence of an object o upon which the action a is performed.

3 Possible worlds and multiple contexts in Scone

Automating common-sense reasoning is a task that requires a sufficiently expressive language, a knowledge base in which to store such a large amount of knowledge, and a set of mechanisms capable of manipulating this knowledge, so as to infer new information. The Scone KB project is an open-source knowledge based system, intended to represent symbolic knowledge about the world as an interconnected network made up of node units and links between them. Its principal strength lies in the way in which search and inference are implemented. Scone adopts a marker-passing algorithm[Fahlman, 2006] devised to be run in the NETL machine[Fahlman, 1979]. Despite the fact that these marker-passing algorithms cannot be compared with general theorem-provers, they are indeed faster, and most of the search and inference operations involved in common-sense reasoning are supported: inheritance of properties, roles, and relations in a multiple-inheritance type hierarchy; default reasoning with exceptions; the detection of type violations; search based on set intersection; and the maintenance of multiple, immediately overlapping world-views in the same knowledge base.

One of the main objectives with which Scone was conceived for was to emulate humans' ability to store and retrieve amounts pieces of knowledge, along with matching and adjusting existing knowledge to similar situations. To this end, the multiple-context mechanism implements an effective means to tackle this objective. The multiple-context mechanism also provides an efficient solution by which to tackle a classical problem of Artificial Intelligence, since it is frame problem.

The great potential of the multiple-context mechanism used by Scone can be better stated by using the example described in [Fahlman, 2006]. Since "Harry Potter World" is quite similar to the real world, a new context, "HPW", could be created as an instance of the real world[1]. Nevertheless, there are differences between these two contexts, such as the fact that in the "HPW" context a broom is a vehicle. This fact can be easily stated in the "HPW" without affecting real world knowledge, in the same way that knowledge of the real world could be cancelled so as to not be considered in the "HPW" context. The way in which Scone handles multiple contexts so as to avoid incongruence problems is by activating one context at a time. By doing this, only the knowledge contained in the active context is considered for the reasoning and inference task.

Unless otherwise stated, the knowledge described in a parent context is inherited by the child context. The context itself is also a node and, like the other the nodes, it stores a set of maker-bits. One of these marker-bits is the context-marker. This bit, when enabled, determines the activation of all the nodes and links that are connected to the active context.

3.1 Actions and events in Scone

Representing actions and events in Scone simply consists of defining two new contexts, one describing the world before the action or event takes place and another that represents the state of the world afterwards. The following example describes a simplified definition of the `move` event.

```
NEW-EVENT move
  :roles
     origin is a place
     destination is a place
     moving-object is a person
  :throughout
     origin differs from destination
  :before
     moving-object is located in origin
  :after
     moving-object is located in destination
```

In accordance with the aforementioned representation of the `move` event, `Lisa moves` can be defined as an individual

[1]In Scone terminology, "general" is the context node that holds knowledge about the real world, and "HPW" would be an individual node, connected by an is-a link to the "general" node.

node of the `move event` for the specific occurrence of Lisa moving from the kitchen to the living room.

```
NEW-EVENT-INDV Lisa moves
the origin of Lisa moves is kitchen
the destination of Lisa moves is living-room
the moving-object of Lisa moves is Lisa
IN-CONTEXT before
STATEMENT-TRUE? Lisa is in living-room
  => No
GET the location of Lisa
  => kitchen
IN-CONTEXT after
STATEMENT-TRUE? Lisa is in living-room
  => Yes
```

Note how in the `before` context Lisa is not yet in the living room but when the active context changes from the before context to the after context, the same question is positively answered.

4 Leveraging common sense in modeling and reasoning about actions and events

The work in [Mueller, 2006] enumerates a list of issues that should be tackled by any attempt made to automate commonsense reasoning. The following subsections analyze these issues from the viewpoint of their representation and support in performing inference and reasoning. Recall that the main focus of the proposed approach is to leverage common sense into action planning in Ambient Intelligence. Hence, the knowledge modeled has been basically restricted to aspects concerning actions and events.

4.1 Time and location

Modeling and reasoning about actions and events should be undeniably associated with a theory of time. Here, the approach proposed to model time adopts the time conceptualization of the Event Calculus[Kowalski and Sergot, 1986], augmented with the multiple-context mechanism. A *context* node can be used to capture the knowledge about the state of the world at a specific *time point* or *time interval*. Regarding space, the work in [Bhatt *et al.*, 2010] also resorts to an approach based on the Event Calculus formalism as a mean to model spatio-temporal abduction for action and change.

Considering that this work is mainly intended for action planning in Ambient Intelligence, the interest in modeling and reasoning about location is focused on providing enhanced location services. Nevertheless, the proposed approach is not exclusive to services, but can also be used to represent any aspect regarding location. Open standards have been used for interoperability purposes[2] Additionally, the work in [Bhatt, 2010] advocates the convenience of enhancing commonsensical reasoning mechanisms with qualitative representation and reasoning techniques to deal with space and location issues [Bhatt *et al.*, 2010].

4.2 Effects of events

As mentioned above, the multiple-context mechanism is the most suitable means of modeling the effects of events. In the simplest scenario, the deffinition of a new `context` suffices to capture the knowledge about the effects of events, or even to capture the indirect effects. Nevertheless, some other scenarios require more elaboration when describing the effects of events.

Sometimes, these effects, rather than being univocally determined by the event occurrence, are subject to the existance of certain conditions. Modeling these context-sensitive effects therefore implies considering the possible worlds that may appear as a result of the event, as determined by the given circumstances. For example, the effect of Lisa picking up an object is that of the object being held by Lisa. If we now consider the scenario of a slippery object, the effect of picking up the object does not necessarily imply that the object is being held since it might be dropped. Depending on how careful Lisa is when she picks up the object, the effect will be of the object being dropped or being held. The means of handling these sorts of effects is to define a new `context` for each different constraint value. Hence, in the case scenario of Lisa and the slippery object, three new `context` nodes hold the descriptions of the possible world. These `context` nodes hang from the parent `after context` node: one of the `contexts` describes the effects of picking up a slippery object without paying special attention; a second `context` describes the effects of picking up a slippery object while paying special attention to not dropping it; finally, the last `context` considers the effect of picking up a normal object.

Nevertheless, the constraints that determine the occurrence of certain effects or others cannot always be known or evaluated. For example, if the level of attention that Lisa pays to picking up the object cannot be assessed, there is no way of foreseeing whether the object will or will not be dropped. The non-determinism of those scenarios creates uncertainty which must also be captured in the action description.

The occurrence of concurrent events also requires a special treatment when coincident events involve cumulative, impossible or cancelling effects. For example, it is not possible to enter two different locations at the same time or, if a door is pulled and pushed at the same time, it remains static.

4.3 Common-sense law of inertia

The "*frame problem*" has been addressed here by means of the multiple context mechanism. Note that the `after context` is a copy of the `before context` which captures those aspects of the world that change as a result of the event occurrence. This property, which makes things continue in the same state, is known as the common-sense law of inertia.

The difficulty involved in dealing with the common-sense law of inertia is that of having to capture and model the knowledge concerning delayed effects or continuous change. As stated above, a delayed effects occurs if the kitchen sink has its plug in and someone turns on the tap: after a while the water will overflow. The common-sense law of inertia is also involved with regard to the water level since it keeps on increasing unless the tap is turned off. Nevertheless, the level does not increase endlessly but rather increases until it reaches the height of the kitchen sink. Afterwards, the water overflows until the water level equals the height of kitchen

[2] Open Geospatial Consortium (OGC). OpenGIS Location Services (OpenLS): Core Services. http://www.opengeospatial.org/standards/ols.

sink.

The event calculus notion of *fluent* is here adopted to deal with these properties that change over time, such as the water level in the open tap example. At each `time instant` the world must be modeled to capture the value of the changing property.

```
NEW-EVENT turn-on faucet
   :roles
      faucet-liquid is a liquid
      faucet-drain is a drain
      faucet-valve is a valve
      level-of-faucet-drain is a FLUENT
    :before
      current-time is T0
      faucet-valve is turned-off
      level-of-faucet-drain is empty
    :after
      faucet-valve is turned-on
      IN-CONTEXT time-instant T1
      level-of-faucet-drain equals (flow * (
          elapsed-time / base-area))
      IN-CONTEXT time-instant T2
      level-of-faucet-drain equals full
      faucet-liquid is dropped-off
```

4.4 Default reasoning and mental states

Default reasoning alludes to the fact that common-sense reasoning is usually performed in uncertainty. For example, the result of turning the fans power switch to on will be that the fan will start spinning around. However, what if the fan is not plugged or it is not working? Most of the time there is no complete information about all these details, so performing default reasoning with exceptions is the most appropriate way in which to handle incompleteness.

In Scone, default reasoning with exceptions is handled by means of *cancel-links*. Please, refer to the work in [Fahlman, 2006] for further information on this subject.

Reasoning about mental states has also been previously addressed. The work in [Chen and Fahlman, 2008] proposes an approach based on " *mental context*" so as to model mental states and their interactions.

5 Action planning with common sense

As has already been mentioned above, the main difficulty faced by systems for Ambient Intelligence lies in coping with innovation. Surveillance contexts typically provide an ideal scenario for unforeseen situations to take place. Furthermore, in most cases, the system will be prompted to elaborate a response in order to manage the unexpected event. A simulated intrusion in a surveyed building poses an interesting scenario in which to asses the performance, regarding action planning, of the proposed model.

First, the presence sensor installed in the servers' room detects an intruder break-in. The guards are automatically notified with the sensor detection. One of the system's goals under these circumstances is to identify and to locate the intruder.

```
IN-CONTEXT intruder-intention
GET the intention of intruder
  => Not known
IN-CONTEXT intruder-break-in
GET the location of intruder
   => servers-room
STATEMENT-TRUE? guards are notified of
    intruder-location
   => Yes
GET the identification of intruder
   => Not known
STATEMENT-TRUE? intrusion alarm status is on
   => Yes
```

The sound of the alarm makes the intruder aware that his presence has been detected. He therefore decides to run away. Meanwhile, the guards are in their way to the servers' room.

```
NEW-EVENT-INDV intruder-leaves-room
intruder is the agent
server-room is the object
IN-CONTEXT intruder intention
GET the intention of intruder
  => Too many
```

After the intruder leaves the room, his location is no longer the servers' room. On the contrary, the intruder is moving through the building in an attempt to escape without being caught. This state of affairs leads to the need for a plan to pursue the goal of locating the intruder. The location of a person is one of those properties that may need to be released from the common-sense law of inertia while the person is moving. Bearing this in mind, the trajectory of a person in movement can be inferred from the successive locations at three consecutive moments in time.

```
SET-FLUENT intruder is located in loc0 at t0
SET-FLUENT intruder is located in loc1 at t1
SET-FLUENT intruder is located in loc2 at t2
STATEMENT-HOLDS? intruder is moving
   => Yes
GET-FLUENT intruder location at t3
   => (covered-distance / elapsed-time) * t3
```

Now, at time instant `t3`, let us say that the intruder's presence cannot be distinguished at the expected location `loc3`. So what has happened? Well, in between `loc2` and `loc3` there is a room. What makes a person abandon the moving trajectory followed?

```
RELEASE-FLUENT intruder location
  => location fluent released
LIST-EVENTS-CAUSING moving-object abandons
    trajectory-of-move
  => enter, stop, sit, jump, lay down, ...
```

Given the plausible events, the system then becomes engaged in proving which of the actions has certainly taken place. The means of verifying this is to check whether the current context is consistent with any of the `after context` of the plausible actions.

```
LIST-AFTER-CONTEXT stop
   => 1. RELEASE-FLUENT moving-object from
         location
      2. the location of moving-object is
         current-location
STATEMENT-TRUE? the location of intruder is
    loc3    => Not known
LIST-EVENTS-REQUIRING the location of thing
    is place => capture, sense, notice, ...
GET service performing capture
```

```
   => video-recording, face detector,
        fingerprint reader, etc.
NEW-INDV shp3 is shape
   the center-of-shape sph3 is loc3
GET video-recording in shape shp3
   => videoRec-at-shp3
LIST-EVENTS-PRECEDING recording upon person
   => focus-person
LIST-EVENTS-PRECEDING focus-person
   => detecting-face, detecting-smile,
        detecting-temperature, etc.
GET service performing detecting-face
   => face-detector
GET face-detector in shape shp3
   => faceDet_at_server
STATEMENT-TRUE? the location of moving-object
      is shp3   => No
```

Each of the possible events causing the intruder to abandon the trajectory will be evaluated recursively[3].

The Planning algorithm proposed in [Santofimia *et al.*, 2010] starts with an empty plan, the Π plan, to be completed with the list of actions, provided by services. This course of actions is intended to emulate the demanded non-feasible action. The course of actions is provided as a set of actions performed on objects, A and O respectively, and the results R of accomplishing such actions. The function *resultOf* refers to the returned value obtained as result of instantiating the a_i action.

6 Conclusions and future works

This work is founded on the conviction that systems for Ambient Intelligence should consider common sense as a constituent element. This work uses action planning, enhanced with common-sense knowledge about actions and events, as the cornerstone of the decision making process.

The main contribution of this work is threefold. First, a model for actions and events in Ambient Intelligence is proposed to characterize the Ambient Intelligence domain knowledge. Second, the model is represented and enhanced to consider the key issues of common-sense reasoning. Third, the proposed strategy for action planning is grounded in multiple-context and possible-world semantics.

This work is an improvement on existing approaches for planning in Ambient Intelligence when devising ad-hoc tailored solutions, on the basis of the available devices and services. Common-sense knowledge is considered throughout the planning, so rather than constraining the planning solution to context knowledge (explicit knowledge), implicit knowledge leads to more appropriate solutions. In the aforementioned case scenario, please note how the trajectory of the intruder has been devised. Also note how the common-sense law of inertia has been used to infer that if the person is not where he was supposed to be, he must have been affected by a particular event. It has been demonstrated above that the stop event is not considered possible, since the current state of the world does not match the after context of the stop action.

[3] http://sites.google.com/site/csrijcai11/

References

[Allen, 1984] James F. Allen. Towards a general theory of action and time. *Artif. Intell.*, 23:123–154, July 1984.

[Bhatt *et al.*, 2010] Mehul Bhatt, Hans Guesgen, and Shyamanta Hazarika, editors. *Spatio-Temporal Dynamics (STeDy 10)*. ECAI Workshop Proceedings., and SFB/TR 8 Spatial Cognition Report Series, August 2010.

[Bhatt, 2010] Mehul Bhatt. Reasoning about space, actions and change: A paradigm for applications of spatial reasoning. In *Qualitative Spatial Representation and Reasoning: Trends and Future Directions*. IGI Global, USA, 2010.

[Chen and Fahlman, 2008] Wei Chen and Scott E. Fahlman. Modeling mental contexts and their interactions. In *AAAI 2008 Fall Symposium on Biologically Inspired Cognitive Architectures, Washington*. 2008.

[Davidson, 1963] Donald Davidson. Actions, reasons, and causes. *The Journal of Philosophy*, 60(23):685–700, 1963.

[Erol *et al.*, 1994] Kutluhan Erol, James Hendler, and Dana S. Nau. HTN planning: Complexity and expressivity. In *In AAAI-94*, 1994.

[Fahlman, 1979] Scott E. Fahlman. *NETL: A System for Representing and Using Real-World Knowledge*. MIT Press, Cambridge, MA, 1979.

[Fahlman, 2006] Scott E. Fahlman. Marker-passing inference in the scone knowledge-base system. In *First International Conference on Knowledge Science, Engineering and Management (KSEM'06)*. Springer-Verlag (Lecture Notes in AI), 2006.

[Georgeff, 1988] Michael P. Georgeff. A theory of action for multiagent planning. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 205–209. Kaufmann, San Mateo, CA, 1988.

[Hommel *et al.*, 2001] Bernhard Hommel, Jochen Musseler, Gisa Aschersleben, and Wolfgang Prinz. The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24:849–878, 2001.

[Kowalski and Sergot, 1986] Robert A. Kowalski and Marek J. Sergot. A logic-based calculus of events. *New Generation Comput.*, 4(1):67–95, 1986.

[Lenat, 1995] Douglas Lenat. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38:33–38, 1995.

[Mueller, 2006] Erik T. Mueller. *Commonsense Reasoning*. Morgan Kaufmann, 2006.

[Santofimia *et al.*, 2010] Maria J. Santofimia, Scott E. Fahlman, Francisco Moya, and Juan Carlos López. A common-sense planning strategy for ambient intelligence. In Rossitza Setchi, Ivan Jordanov, Robert J. Howlett, and Lakhmi C. Jain, editors, *KES (2)*, volume 6277 of *Lecture Notes in Computer Science*, pages 193–202. Springer, 2010.