Actions and Movements

David Israel
Artificial Intelligence Center
SRI International

John Perry
Philosophy Department
Stanford University

Syun Tutiya
Philosophy Department
Chiba University

Abstract

We present an account of action whose main features are that actions are content properties that agents have in virtue of *(i)* the bodily movements they effect and *(ii)* the wider circumstances in which those movements are effected. The account includes definitions of one action being *a way of* doing another, and of performing one action *by* performing another. Although this account is intended to form part of a theory of intelligent action, including the deliberate and intentional actions of human agents or of autonomous robots, in this paper we abstract from the information processing and cognitive factors involved in such actions.

1 Introduction

Action is the goal of planning, for planning is reasoning about what actions to perform, given certain circumstances, in order to achieve a goal. Most actions, perhaps all, involve movements of an agent's body, or more generally of its movable parts. What is the relationship between actions and movements? Are actions movements that are caused in a special way, for example by way of certain aspects of the mental states of their agents—their beliefs, desires, and intentions? This seems wrong, for our notions of actions involve much more than mere bodily movements, however caused. Consider a simple example: moving a block from one location to another. First, we must determine whether we are thinking of one nonrepeatable particular or of the repeatable kind of action: moving a block from one location to another. If the latter, it is clear that different kinds of bodily movement might be involved. Moving a block from one place to another may be a kind of action, but it is not one kind of movement. Moreover, whether we have in mind a single nonrepeatable event or a repeatable kind, it is also clear that moving a block requires more than a bodily movement, however caused. It also requires a block, and (roughly) a path through which the block can be moved. The action of moving a block cannot be characterized solely in terms of the bodily movements of agents that perform the action; on the other hand, bodily movements can be characterized independently of the actions those movements are associated with when performed in given circumstances.

What is the relationship between actions and movements? On our theory, actions are not special kinds of movements, but *content properties of agents* that agents have in virtue of performing those movements in certain circumstances. We proceed to explain this idea and to show how a theory based on it can be used to develop accounts of relationships between actions that are central to a theory of intelligent action.

It will be useful to begin by examining a report of an action. In the course of this brief examination, we shall introduce much of our theoretical vocabulary, and some of our theory.

Consider:

(1) John turned on the light.

We take (1) to describe an *act*: an unrepeatable event in the past. *Actions*, on the other hand, are temporal properties of individuals (agents), in the sense of being predicable of an individual at a time. Turning on a light is an action, as is turning on some particular light **l**. Various people at various times have had the property of turning on a light; the same is true or could be true for the property of turning on **l**. The act described by (1) is a movement of John's that has the property of resulting in **l**'s being turned on. Thus actions are not kinds of acts, but properties of individuals at times.

(1) does not mention movements, nor use the term "result". But there are movements and results involved in John's action, in his having the property of turning on the light then, and we think these movements and results are the keys to developing a theory of movements, actions and the relations among them. When John turned on the light a complex movement of his body occurred. His elbow straightened somewhat; his upper arm rotated forward and upward at the shoulder. His right index finger bent somewhat while the other fingers of his right hand were bent more into a fist (keeping them out of the

*The research reported in this paper has been made possible by a gift from the System Development Foundation and was conducted as part of a coordinated research effort with the Center for the Study of Language and Information, Stanford University.

way), and his body remained stable and his feet stationary. Given his position in front of the light switch, the movement resulted in the switch being flipped to the *on* position, the appropriate circuit closing, and the light going on. John's turning on the light thus consisted in his *effecting* a movement which, given the circumstances of the movement, had certain results.

Actions are grounded in movements. Most of the things we do, we do by producing effects on the objects around us by moving our bodies.[1] We plan what to do in accord with what we know about relations among actions. Intentional action requires *executing* various types of movements, with knowledge of which effects they will have. Which effects are produced is not merely a matter of which types of movements are executed, however. They also depend on the circumstances in which the movement occurs; not only the immediate circumstances, but ones that are quite remote. In the example above these would include, for example, the continued operation of the electrical generating plant that supplied the lights with power. The type of movement John executed will turn on a light in circumstances like the ones he was in—which light depending on which light is connected to the switch he is standing in front of; in other circumstances it might result in someone being tweaked on the snout, or someone being insulted, or a circus dog being commanded to do a somersault.

Reasoning about action, either in planning/practical reasoning or in plan recognition, must ultimately be grounded in reasoning about movements. In designing robots able to act effectively in a wide variety of environments, we must keep in mind that the things the robot can directly control are the movements of its own effectors; to know how to do things in any of a wide variety of circumstances is to know how to move one's body appropriately in those circumstances. For any intelligent agent there will be a repertoire of types of movements that meet two conditions:

(A) The agent can produce movements of these types in a very large range of circumstances.

(B) The agent knows what effects movements of these types will have in (at least some of) the types of circumstances it encounters.

That is, an intelligent agent must grasp at least part of the causal role of the movements it can produce in the environments in which it is likely to find itself. It must associate with the movement types a pattern or relation between types of environments and the effects movements of that type will have in these environments. We call these relationships the *meanings* of types of movements.

The term "meaning" suggests the perspective that underlies our approach. We regard movements as having propositional contents, and in this regard there is a structural similarity to utterances. The propositional content of an utterance—what is expressed—depends on the type of sentence used and the context and wider circumstances in which it occurs. The propositional content of a movement—what it results in—depends on the type of movement and the context and wider circumstances in which it occurs. The term "meaning" has been used for the relation between circumstances and propositional content associated with sentences. Here we extend it to types of movement and types of result.[2]

Our account is based on a good deal of oversimplification and streamlining, both with respect to agents and the language we use to describe them, and our goals are strictly limited. We assume that the agents in question can be viewed as systems with a set of effectors related by an architecture; the types of movement of the whole system are systematically determined by the types of movements available to the effectors and the architecture. These agents *effect* movements; movements are concrete, unrepeatable particulars that belong to various *types*. Acts, then, are movements effected by agents. An agent who effects a movement of a given type is said to *execute* that type. Such *executions*, which play the role of basic actions in our theory, constitute our first category of actions.

When an agent effects a movement of a given type in certain definite circumstances, that movement will have various results: propositions made true by the effects of the movement. We say that the agent *brings about* these results. This is our second category of actions, which we call *accomplishments*.[3] We consider only these two categories of actions in this paper and focus most of our attention on accomplishments.

Actions are mainly of interest to people insofar as they are done purposively, intentionally, vindictively, and the like. The part of our theory that we present here does not touch any of these interesting features of action. The only reason it doesn't apply to a tree falling as well as to a man shooting is that the former is not an act: trees don't *effect* movements; they just move. Nothing in our theory explains the difference between acts and other movements.

Before we can develop an account of the meanings of movements, we need to develop some ideas about movements themselves. We turn to this in the next section, and to the meanings of movements in §3. In §4, we introduce actions and focus especially on those actions (accomplishments) that can be characterized in terms of the result brought about. We also define certain central relations involving actions and movements. In §5, we briefly discuss related work and the last section contains some conclusions and a preview of further research.

[1] We do not here address the issue of purely mental actions. In the interests of simplicity, we shall be dealing only with the kinematics of movements, abstracting completely from considerations of masses and forces. Finally, we shall not consider actions of maintenance and/or prevention.

[2] We do not assume, however, that an act has a unique content. In [Israel and Perry, 1989; Israel and Perry, 1991], we develop a notion of content that allows an event (act, utterance, etc.) to have multiple contents.

[3] The result that is brought about need not be intended; we are using the term "accomplishment" in what might be called its wry sense, according to which one could focus on a quite unintended result of someone's endeavors, and say, "That's quite an accomplishment".

2 Movements

We take movements to be concrete particulars. They belong to various types, and are effected by agents at particular times, in particular places, and in specific circumstances. The results of a movement depend on the type of movement effected, the agent, time, place and circumstances. Our example focuses on a a certain coordinated movement of arm, hand and finger that we call shall "flicking". As we have noted, the same type of movement can be used to do different things by different agents, in different circumstances, at different times. Thus it is natural to associate relations between circumstances and results with types of movements. We call these relations the *meanings* of the movement types.

The movement types with which we associate meanings involve the whole body and all its effectors. Consider again John's flicking in the immediate neighborhood of the light switch. When we said that that type of action in that circumstance would have as a result that the light gets turned on, we didn't really just imagine John flicking. We imagined him flicking while standing still. If he had moved his feet so as to take a full step backward, while his right arm went forwards and upwards, he would not have turned the light on, but would merely have pawed the air ineffectually. If he had moved his feet so as to take a full step forward, he would not have turned the light on, but merely have banged his knuckle against the wall.

When a person flicks, his arm and hand move as a part of an ensemble of movements and non-movements of other bodily parts. The movements that we deal with in this essay are complex movements of the whole body. People often think of partial movements (and movement types), and 'flicking' can be thought of as a label for such a partial type. In so thinking, they are focusing on a (perhaps) complex partial movement and ignoring the movements (or non-movements) of the rest of his body. These latter form the *movement context* for the salient partial movement. The flicking is salient because it constitutes the increment, *given* the movement context and the wider circumstances, necessary for turning on the light. In a more complete account, we would need to be able to keep track of these other movements; to have a theory of movements of persons and other systems, we need to relate it to a theory of the movements of their parts. For our purpose in this paper, however, a very simple conception will suffice.

3 Meanings

Let us suppose that John executed the flicking movement and turned on the light intentionally. We can imagine him realizing that he was standing right in front of and in easy arm's reach of a switch that he believed was connected in the appropriate way to the light. Why does he effect the movement that he does?

We might represent what John knows about movements that explains his doing what he did, as a generalization about (i) movements of the body, (ii) complex types of movement (iii) circumstances in which those movements occur, and (iv) results.

(2) If I effect a movement of the flicking type, while otherwise standing still, in circumstances in which I am standing directly in front of and within easy arm's reach of a light switch of a certain kind that is in on the *off* position and which is correctly connected up to a functioning light, then a result of such a movement will be that the light so connected will get turned on. Moreover, if I simply stand still, then, in those very same circumstances, the light will not get turned on.

The second sentence should not be interpreted as saying that there is no other way, in those circumstances, for John to bring it about that the light is turned on. It simply says that if John's total body movement is a sufficient condition, in the circumstances, of the light being turned on, then his flicking movement, in particular, is a necessary part of that sufficient condition.

Let us now attempt to generalize and abstract, by way of the following generalization over movements (m) involving parameters for total movement types (M), circumstances (C, plus auxiliary parameters for objects and relations involved in C), and results (P, plus auxiliary parameters).

(*) Any movement m that is of type M, that is effected in circumstances of type $C(x_1, \ldots, x_n, m)$, will have as a result that $P(x_i, \ldots, x_l)$ $(1 \leq i \leq l \leq n)$.

Given (*), we can associate a relation between (types of) circumstance, and (types of) results with the movement type M. We call this its *meaning* and denote it as $[\![M]\!]$. Thus we say

$[\![M]\!](C, P)$ iff (*).[4]

4 Actions

We distinguish two categories of actions, *executions* and *accomplishments*.

4.1 Executions

Executions are actions defined simply by the types of movements executed. We use $\mathcal{E}M$ as short for "executes M"

$\mathcal{E}M(\alpha, t, m)$ iff

1. α executes m at t
2. m is of type M

We assume that the results of a movement of type M occurring are identical with that of an agent executing a movement of type M. Someone interested in the theory of dance, for example, might be interested primarily in executions; typically, however, both as agents and theorists, we are not primarily interested in executions, but in accomplishments.

[4] Note that C is not quite a property or type of circumstance, and P is not quite a proposition. They are what might be called parametric properties and propositions.

4.2 Accomplishments

Accomplishments are actions defined by results; they can be reported by way of a certain canonical form: α *brings it about that* P. We use $\mathcal{B}P$ as short for "brings it about that P".[5] Later, in §5.1, we shall introduce a formal language within which to model some of the logic of accomplishments; in this section, as in the previous sections, our treatment is informal.

$\mathcal{B}P_{x_i,\ldots,x_l}(\alpha,t)$ iff
$\exists M, m$, and C such that

1. $\mathcal{E}M(\alpha,t,m)$
2. $C(x_1,\ldots,x_n,m)$
3. $[\![M]\!](C(x_1,\ldots,x_n,m), P(x_i,\ldots,x_l))$.

4.3 The *way of* relation

Our account allows us to analyze various important relations between actions.

Consider the following piece of practical advice or expression of commonsense know-how:

(3) Flipping the light switch to the *on* position is a way of turning on the light (to which the switch is connected in the appropriate way).

What do we mean when we say that flipping the switch to the *on* position is *a way of* turning on the light? There is a relativity to circumstance that is suppressed. We really mean that flipping the switch to the *on* position is a way of turning on the light in certain circumstances C—when the wiring is installed, the fuse is not blown, the power is on, etc. On the view sketched so far, any case (act) in which an agent flips the switch will involve that agent's executing a movement type that, in the given circumstances, has as a result that the switch is flipped. The same is true for any case of an agent's turning on the light. When we say that accomplishing the first is a way of accomplishing the second, we are claiming that however, in those fixed circumstances, you bring about the first, you will have brought about the second.

It will help to introduce the concept of an execution being a *mode of* an accomplishment in a circumstance:

$\mathbf{MO}(C,\mathcal{E}M,\mathcal{B}P)$ iff $[\![M]\!](C,P)$.

Now let C be fixed as above. In C, bringing it about that the switch is flipped to the *on* position is a way of bringing it about that the light is on iff any type of (total body) movement M which is a mode of bringing it about that the switch is flipped to the *on* position is a mode of bringing it about the light (to which the switch is connected) is turned on. More generally, we define a family of two place way of (**WO**) relations between accomplishments parameterized by C:

$\mathbf{WO}_C(\mathcal{B}P,\mathcal{B}Q)$ iff $\forall M \quad \mathbf{MO}(C,\mathcal{E}M,\mathcal{B}P) \Rightarrow \mathbf{MO}(C,\mathcal{E}M,\mathcal{B}Q)$.

[5] We have in mind cases in which the relevant causal chain does not involve the beliefs, desires, and intentions of another agent. We do ultimately intend to accomodate cases, e.g., in which one person brings something about by convincing another to perform do something, but the intuitions we rely on here pertain to the simpler cases.

These relations are pre-orders; they are reflexive and transitive. Reflexivity is a mildly and innocuously counterintuitive property: in C, bringing it about that the switch is flipped is a way of bring it about that switch is flipped. Transitivity is central to means-end reasoning. What of symmetry and antisymmetry? In the circumstances C, bringing it about that the light is turned on is not a way of bringing it about the switch is flipped to the *on* position, but nothing in our account rules out cases in which, relative to some circumstance C', bringing it about that the light is turned on is a way of bringing it about that switch is flipped to the *on* position and vice-versa. Plausible examples of symmetry, though, are hard to come by. We preempt the search for such cases by declaring in advance our readiness to accept the antisymmetry of the **WO** relations, and thus, our acceptance of the claim that, *given a fixed circumstance* C, accomplishments form a partial order.

4.4 The *VO* and *PER* relations

In the foregoing, we defined two temporal properties of agents, (i) that of an agent executing a movement of a type at a time, and (ii) that of an agent bringing it about that P at a time. We also introduced relations among action properties and circumstances: (i) the relation of a movement execution property being a mode of accomplishing a proposition in a type of circumstance and (ii) the relation of one accomplishment property being a way of for another accomplishment property, in a type of circumstance. We now bring these together in an analysis of the *by* relation. We do things by doing other things; that is, we perform some actions by performing others. Our analysis of actions involves two categories: executions and accomplishments. So, too our analysis of the *by* relation involves two subrelations. Consider:

(4) John flipped the switch to the *on* position by flicking.

(5) John turned on the light by flipping the switch to the *on* position.

We offer definitions of two relations involving agents and times, that of an agent bringing it about that P *by* executing a movement of a certain type in a circumstance at a time, and that of an agent, at a time, bringing it about that P *by* bringing it about that Q, in a circumstance. These two relations together comprise our analysis of the *by* relation. We use the notation **VO** (to suggest *in virtue of*) for the first and **PER** for the second.

$\mathbf{VO}(\alpha,t,C,\mathcal{E}M,\mathcal{B}P)$ iff
$\exists m$ such that

1. $\mathcal{E}M(\alpha,t,m)$
2. $C(m)$
3. $\mathbf{MO}(C,\mathcal{E}M,\mathcal{B}P)$

Thus (4) is true iff (roughly):

(4′) John effected a movement of the flicking type, in circumstances such that any movement of that type, in that kind of circumstance, would have as a result that the switch directly in front of which that

movement was effected would be flipped to the *on* position.

PER$(\alpha, t, C, \mathcal{B}P, \mathcal{B}Q)$ iff
$\exists M$ such that

1. **VO**$(\alpha, t, C, \mathcal{E}M, \mathcal{B}P)$ and
2. **WO**$_C(C, \mathcal{B}P, \mathcal{B}Q)$.

Thus (5) is true (roughly) iff:

(5′) John effected a movement of a type and in circumstances such that any movement of that type in that kind of circumstance would have as a result that the switch directly in front of which such a movement was effected would be turned to the *on* position and in that kind of circumstance, any type of movement that had as a result that a switch so related to a movement of that type was flipped to the *on* position, would also have as a result that the light to which that switch was appropriately connected would be turned on.

5 Related work

There has been a great deal of work in AI on planning, and more generally on reasoning about action. McDermott and Allen have developed general theories of events and actions, with special attention to their temporal characteristics [McDermott, 1982; Allen, 1984]. In the situation calculus [McCarthy and Hayes, 1969], actions are treated as functions from states (or situations) to states, where these latter are themselves akin to instantaneous snapshots of possible worlds.

Recent work on intention of Cohen and Levesque [Cohen and Levesque, 1990] presents a more complicated picture. They add action constructors from dynamic logic to a multi-sorted first-order modal language, one sort being of event types, so that standard temporal logic operators are introduced by definition. The main focus of their work, though, is elsewhere: on a theory of the rationally balanced cognitive states of rational agents, and in particular, a theory of commitment and intention. Similar remarks apply to the account of Georgeff and Rao [Rao and Georgeff, 1991]. More closely related to the ideas presented here is Goldman's theory of action, especially the account of action generation [Goldman, 1970]; see also [Pollack, 1986] for an application to the problem of plan recognition.

5.1 The Logic of Accomplishment

In a series of papers, [Segerberg, 1982; Segerberg, 1985] and especially [Segerberg, 1989], Krister Segerberg has attempted to exploit the dynamic logic of programs [Pratt, 1979] to provide a formal account of features of human action. In dynamic logic, one posits a set S of total states and the set S^ω of finite or infinite sequences or *histories* of states. One then associates with each program α its intension, the set $S^\omega_\alpha \subseteq S^\omega$ of executions sequences of α. Segerberg focuses on those (human) actions that can be characterized by way of their results and for which there exists a program or *routine* for accomplishing those results. He introduces an operator δ, to be read as "bringing it about" that takes sentences and yields action (accomplishment) terms. Where Φ is a sentence and for any path p, $p(\#)$ the final element of p, the intension of $[\![\delta\Phi]\!] = \{p \mid \exists\alpha\,(p \in [\![\alpha]\!] \,\&\, \forall p'\,(p' \in [\![\alpha]\!] \Rightarrow p'(\#) \in [\![\Phi]\!]))\}$. These accomplishment terms, rather than program terms (as in standard dynamic logic), are then used to form the modal operators of a multi-modal logic.

Segerberg allows all of the standard operators of regular propositional dynamic logic [Fischer and Ladner, 1979]. We shall present here only the sequential composition (';') fragment of the logic, and we shall substitute our $\mathcal{B}$ for Segerberg's δ. We had investigated this fragment independently of Segerberg. Our intuition was akin to his: our movement types were the analogues of the programs of dynamic logic (Segerberg's routines) and, like Segerberg, we had decided to look at a language in which reference to movements (programs/routines) is suppressed.

The language $\mathcal{L}$ of the logic consists of a language for classical propositional logic, plus closure under the accomplishment operators. These are formed from the special delimiters '[' and ']' enclosing accomplishment terms. The sets of *sentences* and *terms* are defined by simultaneous inductive definition, as follows: Where $\mathcal{P}$ is a denumerable set of sentence letters,

1. Each $P \in \mathcal{P}$ is a sentence.
2. The sentences are closed under a functionally adequate set of Boolean operators.
3. If Φ is sentence, then $\mathcal{B}\Phi$ is a term.
4. If τ_1 and τ_2 are terms, so is $\tau_1;\tau_2$.
5. If τ is a term and Φ, a sentence, then $[\tau]\Phi$ is a sentence.

Sentences of the form $[\mathcal{B}\Phi]\Psi$ are to be read as "After bringing it about that Φ, it is the case that Ψ."[6] The axioms are as follows.[7]

1. any tautology
2. $[\tau](\Phi \wedge \Psi) \equiv ([\tau]\Phi \wedge [\tau]\Psi)$
3. $[\tau_1;\tau_2]\Phi \equiv [\tau_1][\tau_2]\Phi$
4. $[\tau]\mathbf{T}$
5. $[\mathcal{B}\Phi]\Phi$
6. $[\mathcal{B}\Phi]\Psi \rightarrow ([\mathcal{B}\Psi]\Theta \rightarrow [\mathcal{B}\Phi]\Theta)$

The rules are as follows:

1. Modus Ponens
2. External Extensionality: If $\vdash \Phi \equiv \Psi$, then $\vdash [\tau]\Phi \equiv [\tau]\Psi$.
3. Internal Extensionality: If $\vdash \Phi \equiv \Psi$, then $\vdash [\mathcal{B}\Phi]\Theta \equiv [\mathcal{B}\Psi]\Theta$.[8]

[6] In any theory in this language, the standard form encountered would be $\Phi \rightarrow [\mathcal{B}\Psi]\Theta$.

[7] These first four axioms cover the sequential composition fragment of regular propositional dynamic logic.

[8] We accepted external extensionality for the sake of simplicity; but were, and remain, dubious about internal extensionality. Still, given Segerberg's treatment, the logic is congruential, and this is a significant technical advantage.

The semantics of $\mathcal{L}$ is given in terms of frames of the following kind: $\langle S, A, \mathcal{B}, \mathcal{P} \rangle$, where

1. S (the domain of possible total states) is a set.
2. A (the set of actions) is a set of binary relations on S which is
   - closed under relational composition (relative product).
3. $\mathcal{B}$ (the action operator) is a function from P to A.[9]
4. $\mathcal{P}$ (the set of propositions) is a set of subsets of S which is
   - closed under set theoretic union, intersection and complement relative to S; and
   - for each $R \in A$, $\mathcal{P}$ is closed under the operator I_R defined as follows: for all $P \in \mathcal{P}$,
     $$I_R(P) = \{s \mid \forall t(\langle s, t \rangle \in R \Rightarrow t \in P)\}.$$

Segerberg presents soundness and (weak) completeness proofs for a logic that extends full (regular) propositional dynamic logic. The completeness proof also yields a proof of the finite model property and decidability.

6 Conclusions and further research

We have sketched an account of the nature of action whose main features are that actions are properties that agents have in virtue of *(i)* the bodily movements they effect and *(ii)* the wider circumstances in which those movements are effected. Though this account is intended to form part of a theory of intelligent action, including the deliberate and intentional actions of human agents or of autonomous robots, we have abstracted quite completely from the information processing and cognitive factors, including sensory-motor control factors, involved in such actions. Some quite preliminary steps in the direction of a richer theory can be found in [Israel, 1987; Perry, 1986; Israel and Perry, 1989] and more directly in [Israel and Perry, 1991]. Finally, we have borrowed a formal treatment, due to Krister Segerberg, of the logic of bringing it about. Further development in this last direction requires extension to the first order case.

7 Acknowledgments

We would like to thank the members of the **RATAG** (Rational Agency) project at CSLI for stimulating discussion and helpful criticism.

References

[Allen, 1984] J. F. Allen. Towards a General Theory of Action and Time. *Artificial Intelligence*, 23(2):123–154, 1984.

[Cohen and Levesque, 1990] P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(3), 1990.

[Fischer and Ladner, 1979] M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18:194–211, 1979.

[Goldman, 1970] A.I. Goldman. *A Theory of Human Action*. Prentice-Hall, Englewood Cliffs, NJ, 1970.

[Israel and Perry, 1989] D.J. Israel and J.R. Perry. What is information? In P. Hanson, editor, *Information, Language, and Cognition*, Vancouver Studies in Cognitive Science. University of British Columbia Press, 1989.

[Israel and Perry, 1991] D. J. Israel and J. R. Perry. Fodor and Psychological Explanations. In B. Loewer and G. Rey, editors, *Meaning in Mind*. Basil Blackwell, 1991. in press.

[Israel, 1987] D.J. Israel. The role of propositional objects of belief in action. Report 72, CSLI, Stanford, 1987.

[McCarthy and Hayes, 1969] J. McCarthy and P. J. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

[McDermott, 1982] D. McDermott. A temporal logic for reasoning about plans and actions. *Cognitive Science*, 6(1):101–155, January 1982.

[Perry, 1986] J.R. Perry. Circumstantial attitudes and benevolent cognition. In J. Butterfield, editor, *Language, Mind, and Logic*, pages 123–133. Cambridge University Press, Cambridge, 1986.

[Pollack, 1986] M. E. Pollack. Inferring domain plans in question-answering. Technical Report 403, Artificial Intelligence Center, SRI International, Menlo Park, CA, 1986.

[Pratt, 1979] V. Pratt. Models of program logics. In *20th Symposium on foundations of Computer Science*, San Juan, October 1979.

[Rao and Georgeff, 1991] A. S. Rao and M. P. Georgeff. Modeling rational agents within a BDI-architecture. In R. Fikes Allen, J. and E. Sandewall, editors, *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, Cambridge, MA, April 1991. Morgan Kaufmann.

[Segerberg, 1982] K. Segerberg. The logic of deliberate action. *Journal of Philosophical Logic*, 11:233–254, 1982.

[Segerberg, 1985] K. Segerberg. Routines. *Synthese*, 65:185–210, 1985.

[Segerberg, 1989] K. Segerberg. Bringing It About. *Journal of Philosophical Logic*, 18(4):327–347, November 1989.

[9] We use the same notation for both the symbol and the function denoted by it.