Formalizing Affordance

Mark Steedman (steedman@cogsci.ed.ac.uk)
Informatics, University of Edinburgh, 2 Buccleuch Place
Edinburgh EH8 9LW, Scotland UK

Abstract

The idea that to perceive an object is to perceive its affordances—that is, the interactions of the perceiver with the world that the object supports or affords—is attractive from the point of view of theories in cognitive science that emphasize the fundamental role of *actions* in representing an agent's knowledge about the world. However, in this general form, the notion has so far lacked a formal expression. This paper offers a representation for objects in terms of their affordances using Linear Dynamic Event Calculus, a formalism for reasoning about causal relations over events. It argues that a representation of this kind, linking objects to the events which they are characteristically involved in, underlies some universal operations of natural language syntactic and semantic composition that are postulated in Combinatory Categorial Grammar (CCG). These observations imply that the language faculty is more directly related to prelinguistic cognitive apparatus used for planning action than formal theories in either domain have previously seemed to allow.

Introduction

The notion of an *affordance* (Gibson 1966) has in its most basic sense of an invariant supporting perception been extremely helpful in directing attention to non-obvious properties of the sensory array relevant to visual and haptic perception, and motor control (Lee 1980; Turvey 1990). In its more general sense of an interaction with the world that a perceived object mediates (Gibson 1979) it has proved equally attractive to a wide range of theoretical positions that have emphasized the fundamental role of the notion of *action* in human cognition (Norman 1988, 1999). This is the sense in which a door "affords" egress and ingress, a knife affords cutting and scraping, and the like. The attraction of this notion is that it seems to offer a way in which perceptual learning can be linked to the goals and actions upon the environment of the learner, an idea that has been followed up by E. Gibson and Spelke (1983), among others. However, its influence in these domains has been limited by two difficulties.

One has been the controversial idea of "direct perception". This is the idea that the perception that a mailbox " affords letter-mailing to a letter-writing human in a community with a postal system" (Gibson 1979, p.139, citing Gibson 1950) is as directly related to properties of the sensory array as time-to-impact is to characteristics of the optic flow field for a diving gannet. It is certainly hard to believe that the perception of such affordances is "direct" in this strong sense, although recognition of mailboxes, like that of everything else, is undoubtedly mediated *in part* by such Gibsonian invariants of the optic array as relative spatial frequency spectra, and acquisition of the mailbox artefact concept unquestionably depends upon the association of such invariants with affordances in the more general sense. I shall ignore the perceptual aspect of affordances here.

A more serious obstacle to the exploitation of the idea of affordances in this general sense has stemmed from the very fact that many such affordances are actions or events. A formal theory of events in their relation to objects that is applicable to such perceptual categorization and/or conceptual representation of artefact concepts—that is, a theory of what the affordance itself actually is, and how it actually works as a basis for effective action in the world—has been lacking.

The Linear Dynamic Event Calculus

The Linear Dynamic Event Calculus (LDEC) combines the insights of the Event Calculus of Kowalski and Sergot (1986), itself a descendant of the Situation Calculus of McCarthy and Hayes (1969) and the STRIPS planner of Fikes and Nilsson (1971), with the Dynamic and Linear Logics that were developed by Harel (1984), Girard (1987) and others.

Dynamic logics are a form of modal logic in which the $\Box$ and $\Diamond$ modalities are relativized to particular events. For example, if a (possibly nondeterministic) program or command α computes a function F over the integers, then we may write the following:

(1) $n \geq 0 \Rightarrow [\alpha](y = F(n))$

(2) $n \geq 0 \Rightarrow \langle\alpha\rangle(y = F(n))$

The intended meaning of the first of these is "in any situation in which $n \geq 0$, after every execution of α that terminates, $y = F(n)$". That of the second is (dually) that "in any situation in which $n \geq 0$, there is an execution of α that terminates with $y = F(n)$".

We can think of these modalities as defining a logic whose models are Kripke diagrams in which accessibility between possible worlds is represented by events. Such events can be defined as mappings between situations or partially specified possible worlds, defined in terms of conditions on the antecedent which must hold for them to apply (such as that $n \geq 0$ in (1)), and consequences (such as that $y = F(n)$) that hold in the consequent.

The particular dynamic logic that we are dealing with here is one that includes the following dynamic axiom, which says that the operator ; is *sequence*, an operation related to *functional composition* over events, viewed as functions from situations to situations:

(3) $[\alpha][\beta]P \Rightarrow [\alpha;\beta]P$

Using this notation, we can conveniently represent, say, a plan for *getting outside* as the composition of *pushing* a door and then *going through* it, written $push';go\text{-}through'$.

Composition is one of the most primitive *combinators*, or operations combining functions, which Curry and Feys (1958) call **B**. It can be defined by the following equivalence with a lambda term:

(4) $\mathbf{B}\alpha\beta \equiv \lambda s.\alpha(\beta s)$

Plans like $push';go\text{-}through'$ could be written in Curry's notation as $\mathbf{B}push'go\text{-}through'$

Situation/Event Calculi and the Frame Problem

The situation calculi are heir to a problem known in the AI literature as the Frame Problem (McCarthy and Hayes 1969). This problem arises because the way that we structure our knowledge of change in the world is in terms of event-types that can be characterized (mostly) as affecting just a few fluents among a very large collection representing the state of the world. (Fluents are facts or propositions that are subject to change). Naive event representations which map entire situations to entire other situations are therefore representationally redundant and inferentially inefficient. A good representation of affordances must get around this problem.

To avoid the frame problem in both its representational and inferential aspects, we need a new form of logical implication, distinct from the standard or intuitionistic $\Rightarrow$ we have used up till now. We will follow Bibel et al. (1989) and others in using *linear* logical implication $\multimap$ rather than intuitionistic implication $\Rightarrow$ in those rules that change the value of fluents.

For example, we can represent events involving doors in a world (simplified for purposes of exposition) in which there are two places *out* and *in* separated by a door which may be *open* or *shut*, as follows:[1]

(5) a. $shut(x) \multimap [push(y,x)]open(x)$
    b. $open(x) \multimap [push(y,x)]shut(x)$

[1]We follow a logic programming convention that all variables appearing in the consequent are implicitly universally quantified and all *other* variables are implicitly existentialy quantified. Since in the real world doors don't always open when you push them, box must be read as *default* necessity, meaning "usually".

(6) a. $in(y) \multimap [go\text{-}through(y,x)]out(y)$
    b. $out(y) \multimap [go\text{-}through(y,x)]in(y)$

Linear implication has the effect of building into the representation the update effects of actions—that once you apply the rule, the proposition in question is "used up", and cannot take part in any further proofs, while a new fact is added. The formulae therefore say that if something is shut and you push it, it becomes open (and vice versa), and that if you are in and you go through something then you become out (and vice versa).

To interpret linear implication as it is used here in terms of proof theory and proof search, we need to think of possible worlds as states of a single updatable STRIPS database of facts. Rules like (5) and (6) can then be interpreted as (partial) functions over the states in the model that map states to other states by removing facts and adding other facts. Linear implication and the dynamic box operator are here essentially used as a single state-changing operator: you can't have one without the other.

The effect of such systems can be exemplified as follows. If the initial situation is that you are in and the door is shut:

(7) $in(you) \wedge door(d) \wedge shut(d)$

—then the linear rules (5) mean that an attempt to prove the proposition in (8) concerning the state of the door in the situation that results from pushing the door will fail because rule (5a) has removed the fact in question from the database that results from the action $push(you,d)$:[2]

(8) $[push(you,d)]shut(d)$

On the other hand, attempts to prove the following will all succeed, since they are all facts in the database that results from the action $push(you,d)$ in the initial situation (7):

(9) a. $[push(you,d)]open(d)$
    b. $[push(you,d)]door(d)$
    c. $[push(you,d)]in(you)$

The advantage of interpreting linear implication in this way is that it builds the STRIPS treatment of the frame problem (Fikes and Nilsson 1971) into the proof theory, and entirely avoids the need for inferentially cumbersome reified frame axioms of the kind proposed by Kowalski (1979) and many others (see Shanahan 1997).

Using linear implication (or the equivalent rewriting logic devices or state update axioms of Thielscher (1999) and Martí-Oliet and Meseguer (1999)) for STRIPS-like rules makes such frame axioms unnecessary. Instead, they are theorems concerning the linear logic representation.

Even in this extremely simplified world, we need a little more apparatus to represent our knowledge about doors in a way which will allow us to make plans in-

[2]We follow the logic programming convention of negation by failure, according to which a proposition is treated as false if it cannot be positively proved to be true.

volving them. We also need to state preconditions on the actions of pushing and going through. Here ordinary non-linear intuitionistic implication is appropriate:[3]

(10) a. $door(x) \wedge open(x)$
$\Rightarrow possible(go\text{-}through(y,x))$
b. $door(x) \Rightarrow possible(push(y,x))$

These rules say (oversimplifying wildly) that if a thing is a door and is open then it's possible to go through it, and that if a thing is a door then it's possible to push it.

We also need to define the transitive property of the possibility relation, as follows, using the definition (3) of event sequence composition:

(11) $possible(\alpha) \wedge [\alpha]possible(\beta) \Rightarrow possible(\alpha;\beta)$

This says that any situation in which it is possible to α, and in which actually doing α gets you to a situation where it is possible to β, is a situation in which it is possible to α *then* β.

If we regard actions as functions from situations to situations, then this rule defines *function composition* as the basic plan-building operator of the system. Composition is one of the simplest of a small collection of combinators which Curry and Feys (1958) used to define the foundations of the λ-calculus and other applicative systems in which new concepts can be defined in terms of old. Since the knowledge representation that underlies human cognition and human language could hardly be anything *other* than an applicative system of some kind, we should not be surprised to see it turn up as one of the basic operations of planning systems.

This fragment gives us a simple planner in which starting from the world (12) in which you are *in*, and the door is *shut* and stating the goal (13) meaning "find a possible series of actions that will get you *out*," can, given a suitable search control, be made to automatically deliver a constructive proof that one such plan is (14), the composition of *pushing*, and *going through*, the door:

(12) $in(you) \wedge door(d) \wedge shut(d)$

(13) $possible(\alpha) \wedge [\alpha]out(you)$

(14) $\alpha = push(you,d); go\text{-}through(you,d).$

One way to produce this proof, which is suggested as an exercise, is via *backward-chaining* from the goal (13) on the consequents of rules (10) using the transitivity rule (11). The situation that results from executing this plan in the start situation (7) is one in which the following conjunction of facts is directly represented by the database:

(15) $out(you) \wedge door(d) \wedge open(d)$

This calculus is developed further in Steedman 1997, 2002 in application to more ambitious plans, such as the "monkey and bananas" problem, and a number of generalizations of the frame problem, using on a novel analysis of *durative* events extending over intervals of time, which are ignored here.

However, we have said nothing yet about the problem of *search* implicit in searching for and identifying such plans.

Formalizing Affordance using LDEC

Although the example is simplified for purposes of exposition (in particular, with respect to the problem of durativity), it provides the basis for a quite general calculus of events. (See Shanahan (1997), Thielscher 1999, and Steedman (1997, 2000b) for related proposals including discussions of ramification, qualification, delayed action, simultaneity, nondeterminism and other standard problems that such representations have to deal with.)

In fact the representation of actions and events in terms of an association of preconditions and consequences with the core event is a very generally applicable one. If the precondition is a conditional stimulus such as a light, and the consequence is a reward, such as food, while the action concerned is pecking or pressing a bar, then it can be considered as a representation of an *operant* in the cognitive sense of Rescorla and Wagner (1972), itself a notion closely related to that of an affordance.

It also provides the basis for a formalization of the relation between objects and their affordances, of the kind that we need in order to talk about perceptual and cognitive learning in non-linguistic animals and prelinguistic children. For example, the facts in (5) and (6) strike me as a pretty good representation of what my cat knows about the affordances of doors. Of course, the representation is perfectly neutral concerning the invariants that afford the perception of doors in the first place, their relation to bodily properties like the size of the cat's head, and aspects relevant to learning such as motor embedding of the actions of pushing and going through, and so on. It is a representation of what sort of thing it is that is perceived and learned. Nevertheless, the representation could be used to explain the transition she made in her perceptual learning from a stage where doors afforded her (6) (going through for purposes of egress and ingress) but not (5) (pushing to open and close), homing in via a set of superstitious and rapidly extinguishing spurious affordances to a correct affordance (5) supporting the motor plan (14) and its internalization as yet another affordance of doors. The representation also suggests a basis for experimentally investigating precise details of the cat's representation of the affordances of doors. (For example, do they afford her the ingress and egress of other cats?) Many of these experiments have already been done—most notably, by Köhler (1925), in his investigations of tool use and planning in Chimpanzees.

One of Köhler's most thought-provoking observations concerning such planning was the following. A chimpanzee which was perfectly capable of consistently using a tool such as a stick to reach otherwise unattainable objects—one to whom sticks afforded reaching—was unable to enact such a plan unless the stick was ac-

[3]The version of linear logic mixing linear and standard implication is is closely related to "Bunched Implication Logic" (see Pym 2001, which gives an extensive treatment of its semantics and proof theory, including a cut elimination theorem).

tually present in the problem situation. Mere availability of a stick in an adjoining room—even one which the ape had recently explored—was not enough to trigger the relevant knowledge and cause the ape to fetch the stick.

This observation suggests that for non-linguistic animals, including those closest to us in evolutionary terms, access to the affordances of objects is tied to immediate perception of the objects themselves, as Gibson believed. For an animal, this is quite a good way of running your planner. If you don't have much control over your physical environment, it is probably better to look at those plans the situation affords, rather than backward chaining to conditions that there may be no way for you to satisfy, say because of the time of year. This in turn suggests, uncontroversially, that affordances like egress are indexed in such animals by object-concepts like *door*, rather than by end-states like being *out*, and that planning proceeds *reactively* by forward chaining from what is the case, rather than backward chaining from the goal.

We can represent such indexing by first defining actions like *pushing* and *going through* as functions like the following derived from (5) and (6):

(16) a. $push(y,x) \rightsquigarrow \left\{ \begin{array}{l} shut(x) \multimap open(x) \\ open(x) \multimap shut(x) \end{array} \right\}$

b. $go\text{-}through(y,x) \rightsquigarrow \left\{ \begin{array}{l} in(y) \multimap out(y) \\ out(y) \multimap in(y) \end{array} \right\}$

(Here $\rightsquigarrow$ reads as "yields". The linear implication symbol $\multimap$ is overloaded to signify linear mapping of state to state accompanied by deletion and addition of facts. Implication is so closely related to functional mapping, and the functions in question are so closely related to the state update or rewrite axioms of the proof theory that this overloading seems unlikely to cause confusion.)

The set of such functions $Affordances(door)$ constitutes the affordances of doors:

(17) $Affordances(door) = \left\{ \begin{array}{l} push \\ go\text{-}through \end{array} \right\}$

The Gibsonian affordance-based door-schema $door'$ can then in turn be defined as a function mapping doors into (second-order) functions from their affordances like pushing and going-through to their results:

(18) $door' = \lambda x_{door}.\lambda p_{Affordances(door)}.px$

The operation of turning an object of a given type into a function over those functions that apply to objects of that type is another primitive combinator called $\mathbf{T}$ or *type raising*. As in the case of composition (4), the effect of this combinator can be defined by equivalence to the corresponding λ-term:

(19) $\mathbf{T}x \equiv \lambda p.px$

Accordingly, (18) can be rewritten:

(20) $door' = \lambda x_{door}.\mathbf{T}x$

Such a concept of doors is useful for reactive planning, and one can add more affordances to $Affordances(door)$ as one's experience increases. It seems quite likely that this is close to the way cats or at least chimpanzees conceptualize doors.

However, in human terms it is a somewhat stultifying representation, in that it restricts the concept to previously encountered events involving doors that one has somehow stumbled across. One would like to have the advantages in terms of efficiency of planning that thinking of objects in terms of their affordances allows, while also being able envisage novel uses for doors—for example, using one as a table, or as a raft—when circumstances demand it. In other words, one would like to be able to generalize (18) over a wider range of affordances, such as the affordances of natural kinds such as flat movable objects, or of other things that you can push and/or go through. However, there are reasons to think our ability to generalize very far beyond natural kinds and directly experienced affordances is quite limited. (For example, people find considerable difficulty in solving those irritating conundrums which require one to see that a pair of pliers affords the weight for a plumbline, or that the box that thumbtacks are packaged in affords a bracket that can be thumbtacked to the wall to provide a support for a candle.) It seems likely that the basis for such limited generalization is partly perceptual, and partly embedded in our modes of interaction with objects, as Gibson insisted.

Combinatory systems that include both composition and type raising are quite expressive—see Smullyan (1985, 1994) for discussion. They have the character of calculi for rebracketing and permuting terms in expressions. Such calculi are closely related to linear logic itself—see Lambek (1988) for discussion. In this connection it is interesting that the theory of Combinatory Categorial Grammar (CCG, Ades and Steedman 1982, Steedman 2000a) implies that the grammar of all languages involves both type-raising of argument categories and composition of predicates.

Combinatory Grammars

CCG, like other varieties of Categorial Grammar, is a theory in which all linguistic elements are categorized or typed as either functions or basic types, and in which syntactic derivation is achieved by syntactic rules corresponding to directionally and categorially restricted versions of a small number of combinators prominently including composition $\mathbf{B}$ and $\mathbf{T}$. Thus it is a theory that makes language look as if it has been built on a preexisting sytem for planning action in the world, and thereby seem less unique as a cognitive faculty than is usually assumed.

While readers must be directed elsewhere for a full presentation, it may suffice for present purposes to merely note that in CCG elements like verbs are associated with a syntactic "category" which identifies them as *functions*, and specifies the type and directionality of their arguments and the type of their result. For example, a ditransitive verb (DTV) is a function from (indirect object) NPs on the right into transitive verbs (TV)—that is, into functions from (direct object) NPs on the right into

VP:[4]

(21) give := $(VP/NP)/NP$

Such a DTV is a (curried) function that can apply to its arguments to yield VP, as follows:

(22)
$$
\begin{array}{ccc}
\text{give} & \text{Bill} & \text{a biscuit} \\
\overline{(VP/NP)/NP} & \overline{NP} & \overline{NP} \\
\multicolumn{2}{c}{\overline{\qquad VP/NP \qquad}{>}} & \\
\multicolumn{3}{c}{\overline{\qquad\qquad VP \qquad\qquad}{>}}
\end{array}
$$

However, the involvement of further combinatory operations engenders a wide variety of coordination phenomena characteristic of all languages of the world, including English "argument-cluster coordination", "backward gapping" and verb-raising constructions in Germanic languages, and English gapping. The first of these is illustrated by the following analysis, from Dowty (1988):

(23)
$$
\begin{array}{cccccc}
\text{give} & \text{Bill} & \text{a biscuit} & \text{and} & \text{Harry} & \text{an apple} \\
\overline{DTV} & \overline{TV\backslash DTV}^{<\mathbf{T}} & \overline{VP\backslash TV}^{<\mathbf{T}} & \overline{CONJ} & \overline{TV\backslash DTV}^{<\mathbf{T}} & \overline{VP\backslash TV}^{<\mathbf{T}} \\
 & \multicolumn{2}{c}{\overline{\qquad VP\backslash DTV \qquad}{<\mathbf{B}}} & & \multicolumn{2}{c}{\overline{\qquad VP\backslash DTV \qquad}{<\mathbf{B}}} \\
 & \multicolumn{5}{c}{\overline{\qquad\qquad VP\backslash DTV \qquad\qquad}{<\Phi>}} \\
\multicolumn{6}{c}{\overline{\qquad\qquad\qquad VP \qquad\qquad\qquad}{<}}
\end{array}
$$

The type-raising and composition rules, indicated by **T** and **B** repectively, guarantee that the semantics of non standard constituents like *Bill a biscuit* is such as to reduce appropriately with a ditransitive verb like *give*. It is in fact a prediction of the theory that such a construction can exist in English, and its inclusion in the grammar requires no additional mechanism whatsoever.

The earlier papers show that no *other* non-constituent coordinations of dative-accusative NP sequences are allowed in any language with the English verb categories, given the assumptions of CCG. Thus the following are ruled out in principle, rather than by stipulation:

(24) a. *Bill to Sue and introduce Harry to George
b. *Introduce to Sue Bill and to George Harry

Examples like (23) have often been described in terms of very powerful mechanisms of "deletion under identity" of missing elements like the verb *give* in the right conjunct. However, unlike CCG, such proposals fail to explain the observation that such deletions preserve word order, in the sense that in both coordinate and canonical sentences of English, *verbs are to the left of their complements*.

This observation is merely the English specific manifestation of a generalization concerning Universal grammar, due to Ross (1970), who noted that when verbs are "deleted" in this way in languages with other "basic" word orders, such as verb-final (SOV) and verb initial (VSO) languages, they always do so in a way that preserves the canonical left-to-right ordering of verb and argument, thus:[5]

(25) VSO: *SO and VSO VSO and SO
SOV: SO and SOV *SOV and SO

Logical and Neurological Relations between Language and Action

The ubiquitous appearance of composition **B** and type-raising **T** in both affordance-mediated action planning of the most elementary sort on the one hand, and universal grammar on the other, strongly suggests that the language faculty in its syntactic aspect is directly hung onto a more primitive set of prelinguistic operations including these combinators, originally developed for motor planning. This hypothesis has strong implications for the theory of evolution and the child's acquisition of language, for which there is considerable circumstantial evidence from neurological and neuroanatomical observations.

The Linear-Dynamic Event Calculus and related linear and STRIPS-like systems offer a way of representing actions in ways that are useful for planning action. This in turn offers a way of capturing affordances of objects, a notion that is relevant to doing so efficiently, and which is therefore relevant to perceptual categorization and concept learning relevant to tool-use. Two combinatory operations of composition and type-raising play a central role in this process. Those same combinators appear in syntactic guise in natural language, where they provide the basis for an explanatory account of language-specific constructions and cross-linguistic universal generalization, and where a considerable body of evidence from neuroanatomy and child development that has been adduced in support of the Motor Theory suggests that planning and language are closely related. LDEC and CCG make that relation look direct enough to explain the fact that the evolutionary advance in question appears to have been very rapid indeed.

It is interesting to speculate upon what such an evolutionary step might be based. One strong candidate is the attainment of the modal and propositional attitude concepts that are necessary to support a theory of other minds—that is, functions over propositional entities. (We have so far glossed over an important distinction between plans, which compose actions of type *state* $\to$ *state*, and grammar, which composes functions of type *proposition* $\to$ *proposition* or *property* $\to$ *property*.)

It is propositional functions that induce true recursion in both conceptual structures and grammar. There is no evidence that apes entertain such concepts. In particular, the most successful attempts to teach apes to use language, notably those involving ASL and other manipulative languages, show a lack of recursive syntax coupled

[4]We here use the "result leftmost" notation in which a rightward-combining functor over a domain β into a range α are written α/β, while the corresponding leftward-combining functor is written $\alpha\backslash\beta$. (α and β may themselves be function categories.) There is an alternative "result on top" notation, according to which the latter category is written $\beta\backslash\alpha$.

[5]Interestingly, SVO languages like English pattern with verb initial languages in this respect, rather than with verb final. This fact and certain apparent exceptions to Ross's generalization arising in languages with more than one "basic" word order are discussed in Steedman 2000a.

with an almost autistic paucity of conversational initiative. Perhaps it is *only* the lack a theory of mind and the associated propositional attitude concepts that holds apes back from developing human language on the basis of their planning abilities, a suggestion consistent with the views of Tomasello 1999.

Acknowledgments

Thanks to to Ric Alterman, Silvia Gennari, Joyce McDonough, Michael Ramscar, Matthew Stone, and Rich Thomason for comments and advice on a draft. Various stages of the research were funded in part by EPSRC grants GR/M96889 and GR/R02450 and EU (FET) grant MAGICSTER.

References

Ades, A. and Steedman, M. (1982). On the order of words. *Linguistics and Philosophy*, 4:517–558.

Bibel, W., del Cerro, L. F., Fronhfer, B., and Herzig, A. (1989). Plan generation by linear proofs: on semantics. In *German Workshop on Artificial Intelligence - GWAI'89*, volume 216 of *Informatik-Fachberichte*, Berlin. Springer Verlag.

Curry, H. B. and Feys, R. (1958). *Combinatory Logic: Vol. I*. North Holland, Amsterdam.

Dowty, D. (1988). Type-raising, functional composition, and nonconstituent coordination. In Oehrle, R. T., Bach, E., and Wheeler, D., editors, *Categorial Grammars and Natural Language Structures*, pages 153–198. Reidel, Dordrecht.

Fikes, R. and Nilsson, N. (1971). Strips: a new approach to the application of theorem proving to problem solving. *AI Journal*, 2:189–208.

Gibson, E. and Spelke, E. (1983). The development of perception. In Mussen, P., editor, *Handbook of Child Psychology, vol. 3: Cognitive Development*, pages 1–76. Wiley, New York.

Gibson, J. (1950). *The Perception of the Visual World*. Houghton-Mifflin Co., Boston, MA.

Gibson, J. (1966). *The Senses Considered as Perceptual Systems*. Houghton-Mifflin Co., Boston, MA.

Gibson, J. (1979). *The Ecological Approach to Visual Perception*. Houghton-Mifflin Co., Boston, MA.

Girard, J.-Y. (1987). Linear logic. *Theoretical Computer Science*, 50:1–102.

Harel, D. (1984). Dynamic logic. In Gabbay, D. and Guenthner, F., editors, *Handbook of Philosophical Logic*, volume II, pages 497–604. Reidel, Dordrecht.

Köhler, W. (1925). *The Mentality of Apes*. Harcourt Brace and World, New York.

Kowalski, R. (1979). *Logic for Problem Solving*. North Holland, Amsterdam.

Kowalski, R. and Sergot, M. (1986). A logic-based calculus of events. *New Generation Computing*, 4:67–95.

Lambek, J. (1988). Categorial and categorical grammars. In Oehrle, R. T., Bach, E., and Wheeler, D., editors, *Categorial Grammars and Natural Language Structures*, pages 297–317. Reidel, Dordrecht.

Lee, D. (1980). The optic flow field: The foundation of vision. *Philosophical Transactions of the Royal Society, Series B*, 290:169–179.

Martí-Oliet, N. and Meseguer, J. (1999). Action and change in rewriting logic. In Pareschi, R. and Fronhöfer, B., editors, *Dynamic Worlds*, pages 1–53. Kluwer, Dordrecht.

McCarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, B. and Michie, D., editors, *Machine Intelligence*, volume 4, pages 473–502. Edinburgh University Press, WEdinburgh.

Norman, D. (1988). *The Psychology of Everyday Things*. Basic Books, New York.

Norman, D. (1999). Affordance, conventions, and design. *Interactions*, 6:38–43.

Pym, D. (2001). *The Semantics and Proof Theory of the Logic of Bunched Implications*. to appear.

Rescorla, R. and Wagner, A. (1972). A theory of pavlovian conditioning. In Black, A. and Prokasy, W., editors, *Classical Conditioning, II*. Appleton-Century-Crofts, New York.

Ross, J. R. (1970). Gapping and the order of constituents. In Bierwisch, M. and Heidolph, K., editors, *Progress in Linguistics*, pages 249–259. Mouton, The Hague.

Shanahan, M. (1997). *Solving the Frame Problem*. MIT Press, Cambridge.

Smullyan, R. (1985). *To Mock a Mockingbird*. Knopf, New York.

Smullyan, R. (1994). *Diagonalization and Self-Reference*. Clarendon Press, Oxford.

Steedman, M. (1997). Temporality. In van Benthem, J. and ter Meulen, A., editors, *Handbook of Logic and Language*, pages 895–938. North Holland, Amsterdam.

Steedman, M. (2000a). *The Syntactic Process*. MIT Press, Cambridge, MA.

Steedman, M. (2000b). *The Productions of Time*. Ms., University of Edinburgh, `http://www.cogsci.ed.ac.uk/~steedman/`.

Steedman, M. (2002). Plans, affordances, and combinatory grammar. *Linguistics and Philosophy*, 25:(to appear).

Thielscher, M. (1999). From situation calculus to fluent calculus: State update axioms as a solution to the inferential frame problem. *Artificial Intelligence*, 111:277–299.

Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press, Cambridge, MA.

Turvey, M. (1990). Coordination. *American Psychologist*, 45:938–953.