

Preliminary Evidence That Conceptual Structure Can Be Learned by Interacting with an Environment

Paul R. Cohen, Tim Oates and Marc S. Atkin

Computer Science Department, LGRC

University of Massachusetts

Box 34610

Amherst, MA 01003-4610

{cohen,oates,atkin}@cs.umass.edu

Abstract

Recent developments in philosophy, linguistics, developmental psychology and artificial intelligence make it possible to envision a developmental path for an artificial agent, grounded in activity-based sensorimotor representations. This paper describes how Neo, an artificial agent, learns concepts by interacting with its simulated environment. Relatively little prior structure is required to learn fairly accurate representations of objects, activities, locations and other aspects of Neo's experience. We show how classes (categories) can be abstracted from these representations.

Introduction

Our goal is to build a baby, or rather, an artificial agent who lives in a simulated environment and who eventually learns to think like a three year old. A virtual infant, if you like. The abilities we want our agent, which we call Neo (nee "Baby" (Cohen *et al.* 1995)), to develop include learning, planning, language, an organized memory containing structured knowledge, attention, emotion and consciousness. Central to these abilities is a conceptual structure, an ontology, a way to "carve the world at its joints." Given a conceptual structure, we can see a developmental path to more advanced thought, including emotion and consciousness (Cohen *et al.* 1995). As Neo develops, we expect to learn much about the nature of accessible representations, the functions of categories, and the representation and roles of goals. We will elucidate the relationship between activity, attention, memory and learning. We hope to demonstrate that an agent can develop sophisticated knowledge from minimal beginnings. The goal of this paper is more modest, however: to show that Neo can acquire concepts and categories from interacting with its environment.

Categorization, as Lakoff points out, is central to human thought (Lakoff 1984). It is also central to Lakoff and Johnson's challenge to Objectivism (Johnson 1987; Lakoff & Johnson 1980; Lakoff 1984), the dominant view in Western philosophy that there is an objective way to represent the world and reason about it. Influenced by Eleanor Rosch's research on categoriza-

tion, Lakoff and Johnson argue that categories are based less on objective features such as color, size, and shape, than on *interactional* properties and relationships, such as "graspable" and "fits-in-my-mouth," which characterize how an agent interacts with its environment. At the same time, AI researchers such as Agre (Agre 1988), Chapman (Chapman 1991), and Ballard (Ballard, Hayhoe, & Pook) have argued for *deictic* or agent-centered representations. Developmental psychologists such as Gibson (Gibson 1990), Spelke (Spelke 1988), Baillargeon (Baillargeon, Spelke, & Wasserman 1985), and Leslie (Leslie 1988) have found evidence that infants as young as four months know a surprising amount about relationships among actions and objects in the physical environment—force, containment, and the like—although it's unclear whether this knowledge is interactional, in the sense of characterizing how the infant itself interacts with its environment. In any case, Lakoff and Johnson make a convincing case that adult conceptual structures are grounded in primitive interactional knowledge that could very well be acquired by infants. Thus, in contrast with Piaget's theory of developmental *stages* (Ginsberg & Opper 1988), we can now envision a continuous developmental trajectory—for conceptual knowledge, at least—beginning with simple, interactional primitives in infancy and becoming more elaborate through abstraction and metaphorical extension as the agent develops (Lakoff & Johnson 1980). Indeed, Mandler's work, to which we owe much, outlines such a developmental trajectory and the empirical evidence for it (Mandler 1988; 1992). One contribution of the current paper is to show that a simulated infant can learn concepts in roughly the way Mandler suggests real infants learn.

There is a strong temptation to see in the infancy literature evidence of nativism, the idea that infants are born with conceptual structures. Obviously, babies' minds have some structure at birth, but we are anti-nativist, minimalist in our approach. We do not agree that babies must be born with theories of the physical world (Carey & Spelke 1994), and in fact we show that Neo can learn concepts given very little prior structure.

Baby World

Neo is a virtual agent who lives in a simulated environment. Babyworld implements Neo's sensations, mental representations, mental and physical activities, and the behavior of objects and other agents that interact with Neo. Babyworld has two parts: one, which we call Neo, implements everything that Neo does, including learning, moving, looking, crying, and so on. The other part, called StreamsWorld, represents Neo's internal and external environment, and it implements events that happen in and around Neo and in response to Neo's actions. Notably, StreamsWorld represents Neo's sensations of its external environment and also internal states such as hunger. No distinction is made between "inside" and "outside"; Neo must learn it.

Neo senses its environment through a collection of *streams*, which are divided into discrete time steps. In each time step t , a stream σ_i holds a *token* τ ; that is, $\sigma_{i,t} = \tau$. Tokens represent *sensations* or processed percepts. For example, one token is "rattle-shape," and it is placed in the appropriate stream whenever Neo's eyes point at an object that is shaped like a rattle; that is, $\sigma_{\text{rattle-shape},t} = \text{rattle-shape}$. The streams that represent Neo's internal sensations include an affect stream that contains tokens such as happy and sad, a pain stream, a hunger stream, and somatic and haptic streams that are active when Neo moves and grasps.

The Babyworld simulator is simple and probabilistic. For example, Neo gets hungry some time after eating, it cries when it is unhappy or in pain; when Neo cries, Mommy usually visits, unless she is angry at Neo for crying, in which case she stays away. Neo falls asleep intermittently; it can move its arm and head, and grasp several objects, including three rattles, a bottle, a mobile, a bunch of metallic keys, and a knife. The latter causes pain. The rattles make noise when shaken.

How Neo Learns

What does Neo's mind contain? Fluents, mostly. Fluents represent things that don't change, or that change in highly regular, predictable ways. Fluents are a step away from the state-based representations of AI planning research, such as situation calculus, toward script-like representations. Fluents explicitly represent events that have duration; in fact Neo learns the mean and variance of the duration of each fluent. The sound made by a rattle is a fluent, so is the sensation of holding the rattle, and so are the visual sensations of the shape and color of the rattle. Of course, the concept "rattle" has all these components, so somehow, fluents for the color, shape, sound and texture of a rattle must be linked up in a single fluent. Neo accomplishes this, building larger fluents from smaller ones, with two simple learning rules that count cooccurrences. First, if Neo notices that two fluents often start and stop simultaneously, it infers that the fluents are parts of a larger one. This rule learns fluents that represent objects and states. Second, if Neo notices that one fluent

often follows another, it infers that both are parts of a larger fluent that represents an activity. The word "often" in these rules hides a statistical inference that fluents cooccur more frequently than one would expect by chance if they were independent.

The important features of fluents as representations are that they represent states or processes with temporal extent (even objects are represented as things that exist over time), they are composable, and they are learned by counting cooccurrences. Although the simplest fluents represent sensations, it's important to recognize that fluents are not identical with sensations. Sensations are tokens in streams; fluents are representations stored in memory. Streams are the locus of Neo's sensory experience, fluents are the locus of Neo's knowledge. And although Neo's earliest fluents are just copies of its sensations, they soon become aggregated and abstracted.

All the examples in this paper are from a single "run" of Neo, lasting 30,000 time steps (roughly eight hours of Neo's life.) Fluents are learned gradually: It might take hundreds or thousands of time steps to find enough cooccurrences to create a fluent, and composite fluents are obviously learned after their components.

Scopes

Having said Neo combines small fluents into larger ones, we should say where the small fluents come from. The first things Neo learns are not fluents, but rather, pairs of streams in which to look for fluents. These pairs are called *scopes*. A stream σ_i is said to change state at time t , denoted $\Delta(i,t)$, when $\sigma_{i,t-1} \neq \sigma_{i,t}$; that is, σ_i changes state at time t when it contains a different token at time t than it did at time $t-1$. Conversely, $\bar{\Delta}(i,t)$ means the stream doesn't change state: $\sigma_{i,t-1} = \sigma_{i,t}$. Neo learns a *scope*, s_{ij} , when streams σ_i and σ_j change together often. Said differently, Neo learns s_{ij} when the joint event $\Delta(i,t) \& \Delta(j,t)$ occurs frequently relative to the joint events $\Delta(i,t) \& \bar{\Delta}(j,t)$ and $\bar{\Delta}(i,t) \& \Delta(j,t)$. To assess the relative frequencies of these events, Neo uses contingency tables like this one:

	$\Delta(\text{rattle-shape},t)$	$\bar{\Delta}(\text{rattle-shape},t)$	<i>total</i>
$\Delta(\text{rattle-color},t)$	2996	945	3941
$\bar{\Delta}(\text{rattle-color},t)$	826	25232	26058
<i>total</i>	3822	26177	29999

This says that the streams *rattle-shape* and *rattle-color* changed state simultaneously 2996 times, and one changed when the other didn't $945 + 826 = 1771$ times. To assess the strength of association between *rattle-shape* and *rattle-color* we square the frequency in the first cell of the contingency table (2996) and divide by the product of the first row and first column margins (3941 and 3822, respectively). The maximum value for this statistic is 1.0, and for the table above it is $2996^2 / (3941 \times 3822) = .596$.

Neo maintains contingency tables for all pairs of streams. When the measure of association for a table exceeds a threshold, a scope is created. Table 1

shows the top ten scopes learned by Neo in a run of 30,000 timesteps (i.e., the scopes with the highest measures of association), and also the ten worst scopes. Notice that the **sight-shape** and **sight-color** streams are more highly associated than any except **do-sleep** and **sleep**. Other high-ranking scopes are (**sound voice**), which makes sense because changes in the sound stream are often produced by changes in the voice stream; and (**tactile-mouth mouth**), which captures the fact that when Neo starts to mouth (i.e., chew on) an object, it gets tactile sensations in its mouth. The worst scopes represent pairs of streams that are not associated. For example, there is no association between sleeping and eating, and none between moving the arm (**do-arm**) and hunger. Note that many scopes include **do-x** streams; for example, (**do-voice voice**). These are the components of Neo's actions: the **do-voice** part represents the sensation of attempting to use the voice, and the **voice** part represents the sensory feedback from actually using the voice. Sometimes, Neo will experience **do-voice** but not **voice**; for example, Neo might try to make a sound (**do-voice**) but be unable because it has an object in its mouth.

Before Neo learns any scopes, its world is a “blossoming, buzzing confusion” of changing token values in 26 streams. Scopes “chunk” streams into covarying pairs. Without scopes, Neo has to learn fluents by searching for associated token values across all 26 streams. For example, Neo has to consider associations between (**sight-color red**) and, for instance, (**sleep asleep**), (**hunger full**) and (**voice screaming**). With scopes, Neo can limit its search for associations. Suppose Neo learns the scope (**sight-color sight-shape**) but it learns no other scopes relating **sight-color** to any other stream. Then, it should look for associations between (**sight-color red**) and tokens in the (**sight-shape**) stream, but it needn't look for associations in any other streams. Empirically, scopes make an enormous difference in the number and quality of the associations Neo learns. Without scopes, Neo learns many thousands of meaningless associations between token values; with them, Neo learns a few hundreds of associations that correspond to objects and activities in its environment.

Base Fluents

Whereas scopes represent the tendency of streams to change state simultaneously, Neo's smallest fluents, called *base fluents*, represent cooccurring tokens within scopes. Suppose stream σ_i contains a at time $t-1$ and b at time t . Then we say token a stops in stream i at time $t-1$, denoted $\neg(i,a,t-1)$, and token b starts in stream i at time t , denoted $\vdash(i,b,t)$. Now suppose Neo turns its head and its eyes alight on a red rattle. Neo will detect two simultaneous events, $\vdash(\text{sight-color},\text{red},t)$ and $\vdash(\text{sight-shape},\text{rattle-shape},t)$. Sometime later, Neo might look somewhere else, which will gener-

ate two simultaneous stop events, $\neg(\text{sight-color},\text{red},v)$ and $\neg(\text{sight-shape},\text{rattle-shape},v)$. Simultaneous start events and stop events are evidence that a single object—in this case a red rattle—or a single activity, is making its presence felt in two streams. Of course, two *unrelated* events could occur simultaneously in two streams, but this sort of coincidence is less likely than the coincidence of related events.

Neo looks for associations between start and stop events within scopes. For example, having the scope (**sight-color sight-shape**), Neo can try to associate **red** and **rattle-shape**. But if Neo lacks a scope for, say, **sight-color** and **arm-speed**, then it will never try to associate **red** with **fast**. Thus scopes prevent Neo from even considering many meaningless base fluents.

Contingency tables count the cooccurrences of start and stop events, and assess whether start and stop events happen simultaneously significantly often. For example:

	$\vdash(\text{color},\text{red},t)$	$\vdash(\text{color},\overline{\text{red}},t)$	total
$\vdash(\text{shape},\text{rattle},t)$	65	27	92
$\vdash(\text{shape},\overline{\text{rattle}},t)$	237	1931	2168
total	302	1958	2260

Of the 2260 times **sight-color** and **sight-shape** changed together, **rattle-shape** became active 92 times, both **rattle-shape** and **red** became active 65 times, and **red** became active but **rattle-shape** didn't 237 times. The conditional probability of **rattle-shape** starting clearly depends on whether something red or non-red started; these probabilities are $65/92 = .71$ and $27/92 = .29$, respectively. Conversely, the conditional probability of something red starting depends on whether something rattle-shaped started. In short, **red** and **rattle-shape** are associated. The strength of their association can be measured many ways; one was described in the previous section. Here, we use a modified G statistic. Because G is sensitive to sample size, all contingency tables are first scaled to maintain their proportions but have their totals equal 100. To scale the table above, each cell value and marginal total would be divided by 226. Then the G statistic is calculated for the scaled table in the usual way.

Neo accepts a base fluent when its contingency table is significant, as measured by the G statistic. Actually, the table above tells us only that **red** and **rattle-shape** often start simultaneously, we also need to establish that they often end simultaneously. For this, Neo maintains another table like the previous one for the events $\neg(\text{sight-shape},\text{rattle-shape},t)$, $\neg(\text{sight-color},\text{red},t)$, and their complements. When Neo has evidence that **red** and **rattle-shape** both start and stop together in their respective streams, and do so more often than would be expected by chance if they were independent, then it forms the base fluent ((**sight-shape rattle-shape**)(**sight-color red**)).

To summarize the story to this point, Neo learns scopes, or pairs of streams that often change together. As soon as it has learned a scope, Neo can use it to

Table 1: The ten best and ten worst scopes learned by Neo in 30,000 timesteps.

10 Best Scopes	Measure of Association	10 Worst Scopes	Measure of Association
do-sleep sleep	1.0	sleep arm-x-angle	≈ 0
sight-color sight-shape	.596	arm-x-angle arm-y-angle	≈ 0
sound voice	.533	do-sleep arm-x-angle	≈ 0
arm arm-speed	.453	do-sleep eat	≈ 0
do-mouth mouth	.326	sleep eat	≈ 0
do-voice voice	.325	tactile-skin tiredness	≈ 0
tactile-mouth mouth	.315	do-arm hunger	≈ 0
sound do-voice	.276	do-sleep arm-y-angle	≈ 0
tactile-mouth do-mouth	.274	sleep arm-y-angle	≈ 0
do-head head	.254	tactile-mouth hunger	≈ 0

learn base fluents, which are scopes instantiated with particular token values, such as ((sight-shape rattle-shape)(sight-color red)). And as soon as Neo has learned some base fluents, it starts combining them into larger structures called composite fluents.

Composite Fluents

Whereas base fluents represent associations between cooccurring token values in streams, composite fluents represent cooccurring fluents. Neo currently forms two kinds of composite fluents. *Conjunctive* fluents are generated when fluents F_1 and F_2 start together significantly often, and they also end together significantly often. Clearly, conjunctive fluents are like base fluents. However, base fluents combine token values into fluents, whereas conjunctive fluents combine other fluents; and conjunctive fluents are not constrained by scopes. Contingency tables, like those described earlier, tabulate the frequencies of the joint events ($\vdash F_1 \ \& \ \vdash F_2$), ($\vdash F_1 \ \& \ \vdash \overline{F_2}$), ($\vdash \overline{F_1} \ \& \ \vdash F_2$) and ($\vdash \overline{F_1} \ \& \ \vdash \overline{F_2}$). The modified G statistic, described above, tells Neo whether the association between F_1 and F_2 is significant. If so, Neo forms the fluent (**AND** F_1 F_2).

The second kind of composite fluent is formed when one fluent starts in the context of another one. Suppose Neo is holding a rattle, and then it starts to chew on the rattle (called “mouthing”). While it is holding the rattle, the fluent ((tactile-hand wood)(hand close)) is active, and when it starts mouthing, the fluent ((tactile-mouth wood)(do-mouth mouth)) will become active. The latter fluent starts in the context of the former. If this happens significantly often then Neo will form the *context fluent*

```
(CONTEXT
  ((tactile-hand wood)(hand close))
  ((tactile-mouth wood)(do-mouth mouth)))
```

The contingency tables for context fluents are updated in a slightly different way from previous tables. When fluent F_2 starts at time $t + i$, Neo checks to see

whether fluent F_1 is active, and if so, it updates the first cell of the contingency table, ($\vdash F_1, t \ \& \ \vdash F_2, t + i$). If F_2 starts and F_1 isn’t active, then Neo updates the third cell of the table, ($\vdash \overline{F_1}, t \ \& \ \vdash F_2, t + i$). If F_1 is active but F_2 doesn’t start within a window of i time steps, then Neo increments the second cell of the table, ($\vdash F_1, t \ \& \ \vdash \overline{F_2}, t + i$). The modified G statistic tells Neo whether F_2 starts in the context of F_1 more often than would be expected by chance if F_1 and F_2 were independent.

Chains

Context fluents can be chained together to form multi-fluent sequences. Consider the previous context fluent and the following one:

```
(CONTEXT
  ((tactile-mouth none) (voice cry))
  ((tactile-hand wood) (hand close)))
```

These fluents share a common fluent, ((tactile-hand wood) (hand close)), so may be composed into a chain:

```
(CHAIN
  ((tactile-mouth none) (voice cry))
  ((tactile-hand wood) (hand close))
  ((tactile-mouth wood)(do-mouth mouth)))
```

In words, Neo had nothing in its mouth and was crying, then it grabbed something wooden, then it started to mouth something wooden.¹ Now consider another, very similar, chain:

```
(CHAIN
  ((tactile-mouth none) (voice cry))
  ((tactile-hand plastic) (hand close))
  ((tactile-mouth plastic)(do-mouth mouth)))
```

¹Whereas Neo learns all scopes, base fluents and composite fluents incrementally (in a trial of 30,000 time steps, as described earlier), it builds chains with a batch process at the end of the trial. Eventually, chains will be learned incrementally, also.

The only difference between these chains is the object that Neo grabs and mouths: in the first case it is wooden, in the second, plastic. We may form a *class* of things that Neo can grab and mouth. The chains don't say exactly which objects are in the class, but we know they are either wood or plastic, and they are graspable, and mouthable.

Classes

Note that "graspable" and "mouthable" are *interactional* properties (Lakoff & Johnson 1980) that characterize Neo's activities in its environment. Unlike "texture"—wood or plastic—they are fundamentally *subjective*. What's graspable by one agent isn't necessarily graspable by another. Whereas texture is an inherent property of an object, graspable is a property of the object *and the agent* who may try to grasp it. Objective features such as texture have gotten a bad name because they appear inadequate for conceptual activities such as forming classes and judging similarity; for instance, it is difficult, perhaps impossible, to define a category in terms of necessary and sufficient objective features. One is tempted by the conjecture that categories might be defined in terms of necessary and sufficient interactional features, instead. However, we will try to show that categories are best defined in terms of *activities*, and the apparent superiority of interactional features is due to them describing activities better than objective features.

Although Neo learns chains, we are responsible for using these chains to identify features and form classes. This is how we do it: We match up chains that have the same stream names in the same order, creating an *abstract chain* then form classes of the token values. Consider these chains:

```
(CHAIN
  ((do-arm resting) (arm resting))
  ((do-hand close) (hand close))
  ((tactile-mouth wood) (mouth mouthing)))

(CHAIN
  ((do-arm resting) (arm resting))
  ((do-hand close) (hand close))
  ((tactile-mouth plastic) (mouth mouthing)))

(CHAIN
  ((do-arm move-rt) (arm move-rt))
  ((do-hand close) (hand close))
  ((tactile-mouth wood) (mouth mouthing)))

(CHAIN
  ((do-arm move-rt) (arm move-rt))
  ((do-hand close) (hand close))
  ((tactile-mouth plastic) (mouth mouthing)))
```

Looking only at stream names (e.g., *do-arm*, *arm*, *do-hand*, ...) we see that all these chains are instances of this abstract chain: (*do-arm arm*) → (*do-hand*

hand) → (*tactile-mouth mouth*). What we're seeing in the four chains, above, is two activities:

```
resting arm → closing hand → mouthing something
right-moving arm → closing hand → mouthing something
```

and the "something" in each activity is either wood or plastic. It may not be immediately apparent how these activities identify classes of objects, but in fact the objects that can participate in these activities are just those objects that can be grasped, mouthed, and are either wood or plastic. We know, because we built the Neo simulator, that these objects include Neo's rattles and bottles, but not the mobile, Mommy, or Neo's own hand. (Incidentally, if Neo had run for more than 30,000 time steps, it might have learned that its keys can also participate in the abstract chain, above, in which case the instantiated chains would have included a fluent ((*tactile-mouth metallic*)(*mouth mouthing*)).) The point is that the class of objects that can participate in an activity is identified by interactional and objective features—graspable, mouthable, wooden, plastic or metallic.

What Neo Knew and What It Learned

It is sometimes argued that real babies are born with "faculties" for physical and spatial reasoning (Carey & Spelke 1994), language (Pinker 1995), even reasoning about living things (Keil 1994). Nobody believes that much can be learned without constraints from perceptual systems, effectors, and prior mental structures, so the question is not "whether" but "how much." Although we eventually expect to show that Neo can acquire a rich conceptual structure from minimal beginnings, at present, we can show no such thing, because Neo's conceptual structure is quite poor—Neo learned hundreds of fluents and chains but these produced very few classes of objects and activities—and Neo's prior structure is not insignificant. Let us review what Neo was born knowing, and what it learned.

Neo experiences its world through streams, and streams contain tokens such as red and hungry. Neo's learning methods all focus on start (+) and stop (-) events, and thus a notion of events is built in. Neo is born with a method to find correlations among start and stop events in streams. The resulting structures, called scopes, constrain the simplest fluents Neo learns. Neo is born with methods, based on contingency tables, for learning base fluents from scopes and for learning conjunctive and context fluents from both base and composite fluents.

It is further assumed that Neo can maintain contingency tables for all possible scopes, and maintain counts in all its tables. This is a lot of bookkeeping but we believe much of it can be avoided. For example, the *G* statistic requires us to maintain counts for all four cells of contingency tables, but the measure of association introduced in the section on scopes re-

quires only the first cell and the first row and column margins, considerably less bookkeeping.

Now consider what Neo learned: It learned that most of the regularity in its environment takes place in 30 pairs of streams, less than 10% of the $(26 \times 25)/2 = 325$ pairs of streams that it might have focused on. It learned base fluents corresponding to the shape and color of most objects in its environment. It learned the permanent locations of the green mobile (directly overhead) and the crib bars (to the extreme left and right of its field of view). It learned activities, such as grasping an object and mouthing it, or moving its arm and seeing its arm move. It almost learned conditions. For example, it learned a chain that includes ...((do-hand open)(hand open))((tactile-mouth skin)(mouth mouthing)), but it has no way to learn that the first fluent is a condition for the second—that the hand *must* be open to be mouthed. It learned chains from which we abstracted classes that make sense in Neo's environment, such as the class of objects that can be grasped and mouthed, and the class of activities that end in seeing the arm moving fast.

Keep in mind that Neo's actions are largely random: when it grabs an object it *can* mouth it, but it's just as likely to drop it, or move its head. The only structure in Neo's actions is provided by conditions (e.g., it cannot mouth an object it hasn't grasped, and it cannot mouth its hand unless the hand is open) and by a handful of simple behavioral dependencies built into the simulator (e.g., it sometimes grabs what it looks at, and it cries if it gets hungry). Keeping in mind also that Neo ran for only 30,000 time steps, it seems to us that it learned quite a lot.

Conclusion

The goal of the Neo project is to build a virtual infant that learns many of the cognitive skills that we expect from a three-year old. Underlying these skills is a conceptual structure, an ontology, a way to "carve the world at its joints." This conceptual structure identifies classes, and supports judgments of similarity. Following Lakoff, Johnson, Mandler and others, our position is that concepts are based in activities. Neo's fluents represent objects, states, dependencies and activities. We were able to identify classes by examining Neo's learned activities, thus providing the first evidence from this project that conceptual structure can be learned by interacting with the environment.

References

Agre, P. E. 1988. The dynamic structure of everyday life. Technical Report 1085, MIT Artificial Intelligence Laboratory, Cambridge MA.

Baillargeon, R.; Spelke, E.; and Wasserman, S. 1985. Object permanence in five-month-old-infants. *Cognition* 20:191-208.

Ballard, D. H.; Hayhoe, M. M.; and Pook, P. K. Deictic codes for the embodiment of cognition. Computer Science Department, University of Rochester.

Carey, S., and Spelke, E. 1994. Domain-specific knowledge and conceptual change. In Hirschfeld, L. A., and Gelman, S. A., eds., *Mapping the Mind*. Cambridge University Press.

Chapman, D. 1991. *Vision, Instruction and Action*. MIT Press.

Cohen, P. R.; Atkin, M. S.; Oates, T.; and Gregory, D. E. 1995. A representation and learning mechanisms for mental states. In *Working Notes of the Symposium on Representing Mental States and Mechanisms, AAAI Spring Symposium Series*. Also available as University of Massachusetts Computer Science Department Technical Report 95-53.

Gibson, E. J. 1990. *An Odyssey in Learning and Perception*. MIT Press.

Ginsberg, H. P., and Oppen, S. 1988. *Piaget's Theory of Intellectual Development*. Prentice Hall. 3rd Edition.

Johnson, M. 1987. *The Body in the Mind*. University of Chicago Press.

Keil, F. 1994. The birth and nurturance of concepts by domains. In Hirschfeld, L. A., and A. Gelman, S., eds., *Mapping the Mind*. Cambridge University Press.

Lakoff, G., and Johnson, M. 1980. *Metaphors We Live By*. University of Chicago Press.

Lakoff, G. 1984. *Women, Fire, and Dangerous Things*. University of Chicago Press.

Leslie, A. M. 1988. The necessity of illusion: perception and thought in infancy. In Weiskrantz, L., ed., *Thought Without Language*. Oxford University Press (Clarendon).

Mandler, J. M. 1988. How to build a baby: On the development of an accessible representational system. *Cognitive Development* 3:113-136.

Mandler, J. M. 1992. How to build a baby: Ii. conceptual primitives. *Psychological Review* 99(4):587-604.

Pinker, S. 1995. *The Language Instinct*. Harper Collins.

Spelke, E. S. 1988. The origins of physical knowledge. In Weiskrantz, L., ed., *Thought Without Language*. Oxford University Press (Clarendon).