Object knowledge in infancy: current controversies and approaches

Denis Mareschal

Studies relying on looking-time measures have found evidence of a far more precocious understanding of hidden objects than Piaget originally described. However, there is now a heated controversy surrounding the results from looking-time studies – do they constitute any evidence of a conceptual or explicit understanding of objects? Moreover, even within the looking-time paradigm, young infants show rapid changes in their understanding of what constitutes a legitimate occlusion event, and in their ability to use feature information to individuate or keep track of the number of hidden objects. The picture that emerges from these studies is that young infants have a limited and sometimes fragmented understanding of hidden objects. We suggest that computational modelling could help provide a coherent account of the emergence of object-directed behaviours in infancy, although the fit between current models and existing data remains poor.

No psychologist has had more impact on our understanding of child development than Jean Piaget. Although the popularity of his theories may have waned since his first influential publications in the early 1920s, his work has set the research agenda in cognitive development for well over 50 years. One of the most important questions he raised was how and when infants understood that hidden objects continued to exist beyond direct perception[1,2]. The ability to reason about hidden objects liberates infants from the tyranny of direct perception. According to Piaget, knowledge of hidden objects develops through six stages that span all of infancy, and does not reach adult levels of maturity until the age of two. Piaget based this conclusion on search tasks in which infants were required to retrieve objects that had moved out of sight. The retrieval mode varied with age and depended on the motor maturity of the infant. While originally greeted with scepticism, these early findings were replicated and extended in numerous other studies[3].

Piaget argued that infant object knowledge was closely tied to the sensorimotor system. Infants under the age of two were not endowed with any ability to reason symbolically about hidden (or visible) objects. They acquired the understanding that hidden objects continue to exist through active exploration of the world. What constitutes evidence of an 'understanding' of hidden objects, and whether this knowledge is directly available to infants or acquired through experience, remain two of the most controversial questions of contemporary infancy research.

The past 15 years have seen a veritable explosion of studies suggesting a far more precocious understanding of hidden objects by young infants. These studies rely on the finding that infants direct more attention (look longer) at novel or unexpected stimuli. Two principal methods are used:

(1) In the habituation–dishabituation paradigm, infants are shown an event involving objects, until their attention to the event drops below a set criterion. They are then shown related events that differ from the original along some object-relevant dimension. One event is consistent with the belief that hidden objects continue to exist during occlusion and the other event is consistent with the converse belief. How infants transfer what they learned in the familiarization event to the test events allows investigators to evaluate what the infants believe about hidden objects.

(2) In the violation-of-expectation procedure, infants are initially shown a familiarization event. They are then shown a related impossible event (in which some physical property of a hidden object has been violated) or a related possible event (in which no physical property of the hidden object has been violated). Longer looking at the impossible than at the possible event is interpreted as suggesting that the infants' knowledge of hidden objects lead them to perceive the impossible event as more novel or unexpected.

Denis Mareschal is at the Centre for Brain and Cognitive Development, Birkbeck College, Malet Street, London, UK WC1E 7HX.

tel: +44 20 7631 6582
fax: +44 02 7631 6312
e-mail: d.maresschal@bbk.ac.uk

 PII: S1364-6613(00)01536-9

In a seminal series of papers, Baillargeon, Spelke and Wasserman[4,5] used the latter methodology to show that infants as young as 3.5 months understood that hidden objects continued to exist, and understood that the hidden object maintained the physical property of solidity. Further work[6] tracking the development of this knowledge showed that infants under the age of 8 months represented the height and location of hidden objects, understood that a soft hidden object could be compressed, and could even recognize an appropriate series of actions to retrieve a hidden object. Finally, by 13 months infants could use the size of a protuberance to infer the size of an object under a cloth. Habituation–dishabituation studies revealed that 4-month-olds understood that rolling or falling objects, which were temporarily occluded, moved continuously in time along an unobstructed trajectory[7], and that by 8 or 10 months they could also rely on inertial constraints to the same effect[8]. Taken together, this stunningly sophisticated understanding of hidden objects in very young infants has been used as evidence that the core of object knowledge is directly available to infants from birth and that learning plays only a small role in the understanding of hidden objects[9].

These results stand in sharp contrast to Piaget's finding that infants will not reach around an occluding screen to retrieve a hidden toy until 9 months[1,2]. There appears to be a developmental lag between infants' knowledge of hidden objects (as measured by looking times) and their ability to demonstrate that knowledge in active retrieval tasks (such as the manual search used by Piaget). This lag is not simply due to motor co-ordination problems as, by 6 months, infants have sufficient manual dexterity to remove the occluding screen or to reach around it[10]. Several explanations of this lag have been proposed. One possibility is that there are distinct and encapsulated perception and action domains. Moreover, knowledge in these two domains develops at different rates[8,9]. A second possibility is that object representations are developing continuously and that the representation required to elicit a look response is simply a precursor of that required to elicit a reach response[11,12]. Finally, it may just be that infants lack the planning or problem-solving abilities required to solve the manual retrieval tasks[6,13]. This last possibility was recently challenged when infants were found to solve retrieval problems requiring similar or greater planning abilities than in the traditional Piagetian tasks[12].

These initial findings begin to suggest that early infant knowledge of hidden objects may not be as robust as initially believed. The knowledge seems to be modality specific and task dependent. Infant performance depends on the response requirements of a task (e.g. general inspection of an event versus active search) as well as the response modality (e.g. looking versus reaching)[14,15].

Work over the past three or four years has continued to suggest that knowledge of hidden objects is limited. This conclusion is based on three new independent lines of research: (1) investigations into rapid occlusion-related learning in early infancy; (2) the discovery of surprising lacuna in infants' understanding of hidden objects; and (3) a renewed intensity in the questioning of whether perceptually based responses constitute evidence for an understanding of hidden objects. The first two lines of research show that differences in apparent knowledge occur even when the same response modality is used to test that knowledge. The final line of research re-examines what might constitute evidence of object knowledge and questions whether looking-time measures reveal anything about object representations. In the rest of this article I will briefly present developments along these three lines of research before suggesting a methodological approach familiar to cognitive scientists that has the promise of offering a synthesis over the diverse range of infant behaviours.

When is an occlusion an occlusion?

Over the past 5 years, researchers have started to explore how much even younger infants than those studied previously understand about hidden objects. Although infants older than 3.5 months appear to appreciate that occluded objects continue to exist, recent studies suggest that the ability to predict when an object should be completely hidden is poor at younger ages and undergoes systematic development. Building on prior research looking at infants' understanding of occlusion events[16,17], researchers explored when infants would expect to see a translating object appear through a gap in an occluding screen[18,19]. They presented infants with an event in which a toy mouse moved across a stage, temporarily passing behind an occluding screen (Fig. 1). In some cases, the middle segment of the occluding

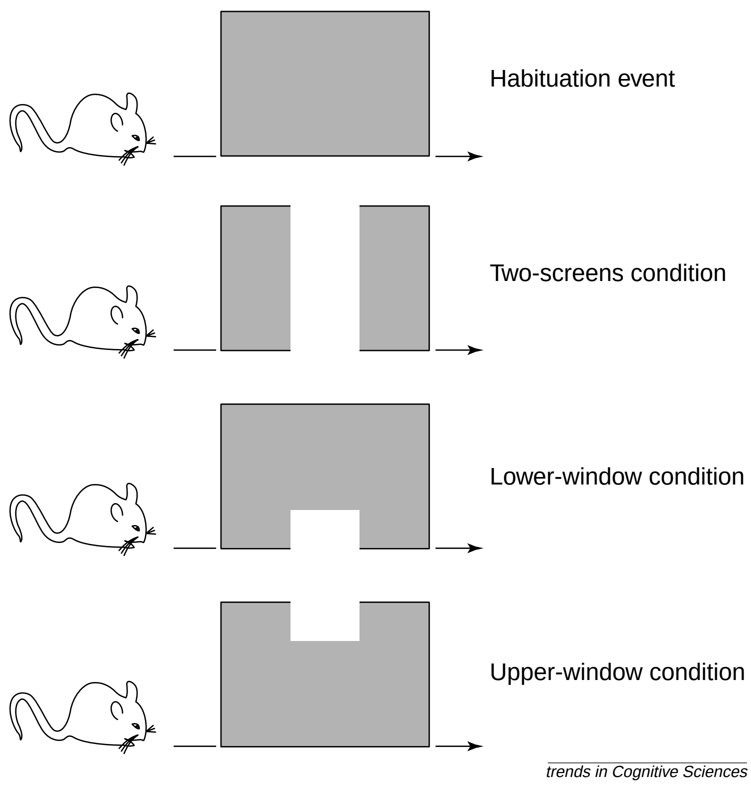

Fig. 1. Example of the habituation and test conditions in Refs 18,19. Infants were presented with events in which a mouse moved across a stage and passed temporarily behind an occluding screen. In some cases, the middle segment of the occluding screen was removed (**top**; the two-screen condition), in others the lower part of the middle segment of the screen was removed (**middle**; the lower-window condition) and in others the top of the middle segment of the screen was removed (**bottom**; the upper-window condition). The mouse was small enough to appear in the lower-window condition and not large enough to appear in the upper-window condition.

screen was removed, leaving a complete gap between what appears as two distinct narrow screens (the two-screen condition). In other cases, either the top half of the middle segment of the screen was removed (the upper-window condition) or the lower part of the middle segment of the screen was removed (the lower-window condition). The mouse was small enough to allow it to appear through the window in the low-window condition but not in the upper-window condition. The authors found that 2.5-month-olds were surprised when the mouse failed to appear between the screens in the two-screen condition, but not when the mouse failed to appear in the lower- or upper-window conditions as it moved across the stage and behind the occluding screen. This was interpreted as evidence that 2.5-month-olds have only a crude concept of occlusion whereby objects are either behind or not-behind an occluder, and hence expect the object to be either not-visible or visible. Three-month-olds expect objects to remain hidden when passing behind an occluder with a continuous lower edge, and to become visible when passing behind an occluder with a discontinuous lower edge. Finally, 3.5-month-olds actually attend to the height of the object in comparison to the height of the occluder, and expect a tall object to become visible when passing behind a short screen. However, learning is not complete by 3.5 months. Infants still have to learn that the width of an object relative to the screen will predict whether it is fully hidden[20], that the speed of an object determines how soon it will appear from behind the occluder[8] and that the speed and width of the occluder determine how long it will take to pass by an aperture in the screen[21].

Even after 3.5 months, infants' understanding of occlusion events is fragile. For example, they seem to have great difficulty in predicting when an object lowered into a container should be completely hidden[18,19]. In these studies 4.5- and 7.5-month-old infants saw two events in which an object was lowered into a cylindrical container or lowered behind a convex screen of equal dimensions to the container. The younger infants failed to appreciate that the height of the object being lowered would predict whether it was completely occluded in the containment event, but succeeded to do so in the occlusion event. So, even though 3.5-month-olds appear to reason about the height of a hidden object in occlusion events[22], it is not until 7.5 months that they can use height to reason about containment.

Taken together, these findings suggest that from 2.5 to 7.5 months of age, what young infants 'know' about events in which an object disappears is still developing, but continues to be closely tied to a particular kind of event. Thus, young infants appear not to have a consistent understanding of hidden objects. Even when the same response modality is used to assess knowledge (i.e. looking), different tasks lead to a demonstration of different levels of understanding at a given age (e.g. 4-month-olds' understanding of occlusion and containment tasks) and different rates of learning.

The 'what' and 'where' of hidden objects

Infants appear very good at using spatial or temporal information to keep track of the number of hidden objects behind a screen. For example, in the two-screen condition described above, even 2.5-month-olds will expect to see two objects if the screens are removed[18,23,24]. However, this ability is sometimes coupled with a remarkable limitation when a single large screen is used to occlude the possible objects. When watching an event that involves objects being placed behind an occluding screen, or brought out from behind the screen, 5.5-month-olds can keep track of the number of objects behind the screen when that number remains between one and three[25,26]. However, they show no surprise if the screen is removed to reveal the correct number of objects, but objects with completely different surface features[27]. These infants appear to keep track of the objects' spatial and temporal properties (location) but not their surface features.

Further evidence also initially suggested that even 10-month-olds base their individuation (enumeration) of hidden objects on spatial–temporal information only. If shown a toy duck and a toy car appearing and disappearing, one at a time, from behind an occluding screen, 10-month-olds will expect a single object when the screen is removed. In contrast, if the duck and the car are shown appearing and disappearing simultaneously, the 10-month-olds will expect two objects when the screen is removed[28]. Not until 12 months did infants appear to base their individuation of hidden objects on feature information in this task. However, 12-month-olds still fail to use colour as a cue to individuate objects, even though they show evidence of remembering the colour of the objects[29].

However, here again, there is evidence that small changes in task requirements lead to radical changes in apparent knowledge. A distinction has recently been drawn between event-mapping tasks, in which infants watch two events and judge whether the two are consistent (e.g. the work on individuation above), and event-monitoring tasks, in which infants see a single event and judge whether successive portions of the event are consistent. Infants seem to perform better at an earlier age on the latter type of task. In the event-monitoring task, infants are shown two types of occlusion events and their looking times at these two events are compared. No transferal of knowledge from one context (occlusion) to another context (no occlusion) is required. When tested with event-monitoring tasks 4.5-month-olds can use shape and size information, 7.5-month-olds can use texture pattern information, and 11.5-month-olds can use colour information to individuate objects[20,30,31]. Why infant performance differs on these two tasks is a topic of continuing research[32,33].

These findings continue to argue against the notion that infants possess a robust understanding of hidden objects. Young infants are able to keep track of the number of hidden objects, but appear not to keep track of the objects' identity at the same time. Moreover, infant abilities to individuate objects are not consistent across tasks even when the same modality is used to assess their knowledge (i.e. looking). Whether using an object-mapping task or an object-monitoring task, the age at which individuation of occluded objects is demonstrated will differ. However, infants always begin by using spatial–temporal cues for individuation (e.g. shape, size, location) before using surface feature cues for individuation (e.g. colour).

Is seeing the same as believing?

Recently, some researchers have argued that extreme caution should be taken when interpreting studies whose conclusions are based on measuring looking times. Apart from some difficulties in replicating a few of the findings[34], many authors do not accept that expectation-based tasks reveal anything about understanding. These authors argue that a distinction should be drawn between perception and cognition in studies of infant knowledge[35]. The results described above are all based on methods originally designed to assess infant perceptual development, and perception is not the same as knowing[36]. It is debatable whether infants are actually expecting (forecasting) in these object–event studies. Indeed, only after the event has occurred can there be a mismatch with similar events retrieved from memory[37]. Studies that have systematically varied perceptual variables (such as the zone of infant tracking and the amount of time the object is visible) suggest that some of the behaviours attributed to an understanding of hidden objects can be explained in terms of these perceptual variables[38,39]. Moreover, the expectation-based studies often ignore the fact that anticipation itself develops with age[40].

A consequence of this position is to question what kind of 'knowledge' underlies infants' performance on looking-time tasks. Simple perceptual memory accounts have been proposed,[41] with counterarguments also advanced[42]. The crux of this debate is whether infants are guided by a conceptual representation of objects or whether their behaviour is guided by perceptual constraints on information processing, and how this information will be used[43]. In short, the very notion of what constitutes 'knowledge' of hidden objects is itself incoherent.

Steps towards a resolution: process models of object-directed behaviours

Attempts to resolve these questions abound. They range from the more nativist perspective that attributes core reasoning principles to the newborn infant[9,42], to the idea that infants develop context-specific concepts of hidden objects that are refined through the identification of relevant task variables[19]. It has also been suggested that knowledge of hidden objects could be acquired simply through associations and feedback from the environment[12,15,44].

Cognitive science provides a glimpse of how these issues might be resolved. The job of developmental psychologists is to elucidate the processes – the causal chains – by which new behaviours emerge[45]. As long as there is no functional account of how representations are used, that term remains vacuous. By implementing computational models (as is common practice in cognitive science), researchers are forced to be explicit about the nature, scope and power of the knowledge representations they posit, as well as being explicit about processes that operate over those representations[46,47].

There are relatively few computational models of infant object-directed behaviours. Early models took a strong cognitivist stance on behaviour and were thus implemented in rule-based production systems[48–50]. These models often implemented a variant of Bower's Identity Theory of object development (Box 1). Unfortunately, they were basically competence models that described infant behaviours but did not provide a mechanistic account of development. They proposed different sets of rules to describe behaviour at different ages but did not explain how new rules could be acquired or how one set of rules was transformed into another

Box 1. Identity theory

The original Identity Theory (Ref. a) suggested that, rather than studying object permanence, we should be studying the development of object identity. Bower believed that young infants understand that objects continue to exist but that they have difficulty keeping track of objects. Young infants generate a large number of separate object representations for what adults would encode as a single object. Identity Theory has been implemented in a series of production system models (Refs b,c). The models assume that infant behaviours can be ascribed to the use of five action rules subsumed under three conceptual rules:

Rule 1: An object is a bounded volume of space in a particular place *or* on a particular path of movement.

Rule 2: An object is a bounded volume of space of a certain size, shape and colour, which can move from place to place among trajectories. (Note that this rule now integrates feature information with spatial–temporal information.)

Rule 3: Two or more objects cannot be in the same place *or* on the same path of movement simultaneously *unless* they share a common boundary.

The level 1 rule tests for the location of the object, whether this is the expected position based on the previous snapshot, whether the object has volume and whether it is occluded. The level 2 rule adds a test for the object's features before testing for intact boundaries. Finally, the level 3 rule adds another test of features after testing for intact boundaries.

The level 1 models do not rely on feature information to identify an object. Hence, they set up a new object every time there is a change in spatial–temporal information. The level 2 models set up a new object token only when there are changes in the feature representation. Finally, the level 3 models do not set up a new object representation when two objects are contiguous.

In its most recent incarnation, Identity Theory has been proposed as an account of how young infants determine the identity of inanimate objects, as well as of people (Ref. d). In this view, the early representation of objects is not static but dynamic. Infants strive to maintain coherence between their representations and the perceived world. People play an important role in learning about objects because they change features often. Early concepts are used to interpret the behaviour of people and things and are revised in light of later experience (Ref. e).

References

a Bower, T.G.R. (1982) *Development in Infancy* (2nd edn), W.H. Freeman

b Luger, G.F. *et al*. (1983) A model of the development of the early object concept. *Perception* 12, 21–34

c Luger, G.F. *et al*. (1984) Modeling the stages of the identity theory of object concept development in infancy. *Perception* 13, 97–115

d Meltzoff, A.N. and Moore, M.K. (1998) *Infant Behav. Dev.* 21, 201–235

e Gopnik, A. and Meltzoff, A.N. (1997) *Words, Thoughts, and Theories*, MIT Press

Box 2. Attention-based models

Attention-based models of infant object reasoning draw on ideas from adult visual attention (Refs a,b). This approach is related to the FINST (Refs c,d) and OBJECT-FILE models (Ref. e). Attention can span more than one object at a time. Object indexing plays a role in individuation and identification of distinct objects. An index is a mental token that functions like a pointer. It does not contain any feature information and is resource-limited (i.e. there are a small number of them – no more than four). Indices are assigned by location to an object; there is a distinct index for every object. Indices remain with objects over spatial transformations and objects must occupy distinct locations to be assigned different indices. Because there is a limited number of indices, they are reused. If location information is ambiguous, then indices can be assigned by property information.

A related model, INFANT (Refs f,g), implements an attention-limited indexing process within the ACT-R architecture (Ref. h). The model reproduces infant behaviours in early numerical competence studies and individuation tasks without any recognizable numerical representations. Hence, it demonstrates that innate numerical competence is not necessary to explain the behaviour of young infants in these tasks.

References

a Leslie, A.M. *et al.* (1998) Indexing and he object concept: developing 'what' and 'where' systems. *Trends Cognit. Sci.* 2, 10–18

b Scholl, B.J. and Leslie, A.M. (1999) Explaining the infant's object concept: beyond the perception/cognition dichotomy. In *What is Cognitive Science?* (Lepore, E. and Pylyshyn, Z., eds), pp. 26–73, Blackwell

c Pylyshyn, Z. (1994) Some primitive mechanisms of spatial attention. *Cognition* 51, 363–384

d Scholl, B. and Pylyshyn, Z. (1999) Tracking multiple objects through occlusion: clues to visual objecthood. *Cognit. Psychol.* 38, 259–290

e Kahneman, D. *et al.* (1982) The reviewing of object files: object-specific integration of information. *Cognit. Psychol.* 24, 175–219

f Simon, T.J. (1997) Reconceptualizing the origins of number knowledge: a 'non-numerical' account. *Cognit. Dev.* 12, 349–372

g Simon, T.J. (1998) Computational evidence for the foundations of numerical competence. *Dev. Sci.* 1, 71–78

h Anderson, J.R. (1993) *Rules of the Mind*, Lawrence Erlbaum

set of rules. More recent (cognitivist) models have turned to attention-based accounts of object processing (Box 2) in an attempt to explain infant behaviours. Unfortunately, these models still fail (by and large) to implement any account of how development might occur.

One mechanistic learning model has implemented a parallel processing version of Piaget's sensorimotor theory of infant development. It tried to show how the co-ordination of intra- and intermodal perceptual motor schemas could lead to a single unified representation of object[51]. Perceptual motor schemas were encoded as 'context–action–result' rules and implemented in a parallel processing machine. Learning consisted of using marginal probabilities to fill in context and results slots in appropriate perceptual–motor schemas. Although this system developed an intricate network of intra- and intermodal schemas, it did not develop according to the pattern described by Piaget.

This initial parallel processing approach is reflected in recent connectionist models of infant object-directed behaviours. These latter models have taken the developmental lag between infants' looking time and reaching behaviours as their starting point.

In one set of connectionist models, a network learns to predict the reappearance of a stationary object from behind a moving screen that temporarily occludes the object (Box 3). Network performance is measured by taking the difference in response of the nodes coding the location of the hidden object when an object should be revealed and subtracting it from the response of the node when an object should not be revealed. An increase in this difference is interpreted as increased knowledge of hidden objects. What this model shows is that object representations that guide the response to objects can be graded[52] and arise though interactions within an environment. No *a priori* object representations are required.

Another connectionist model (Box 4) is more closely tied to the neuropsychological finding that visual object information is processed down two separate routes[53,54]. This model uses a combination of modules to implement dual route processing. One route learns to process spatial–temporal information while the other route learns to process feature information. Finally, a response module recruits and coordinates the representations developed by the other modules as and when required by a response task. The route specializations emerge as a result of the different associative mechanism in each module.

This latter model suggests that the developmental lag arises from the added computational demands of integrating information across separate cortical representations in an active, voluntary retrieval task. Infant responses in the expectation-based tasks are environmentally driven and can thus be viewed as involuntary responses[55]. When there is a mismatch between the information arriving from the environment and the information stored within the object representations, this mismatch is manifested by a surprise response. Such responses can arise by the independent violation of information contained in either of the two cortically separable processing streams. Closer inspection of studies claiming that young infants remember features of hidden objects reveals that these are always spatial or temporal features (e.g. location, size, velocity, coarse shape). These features are all processed by the dorsal route[53]. In contrast, selective manual retrieval responses are voluntary[55]. Infants do not reach equally for all objects. Moreover, to make a voluntary retrieval response infants require access both to the ventral representation (to identify the object) and to the dorsal representation (to localize the object).

Voluntary retrieval responses are not the only responses that can require integration between pathways. Surprise responses (manifested by increased looking times) may also require integration. However, because it is the requirement to integrate information across cortical

Box 3. Connectionist models

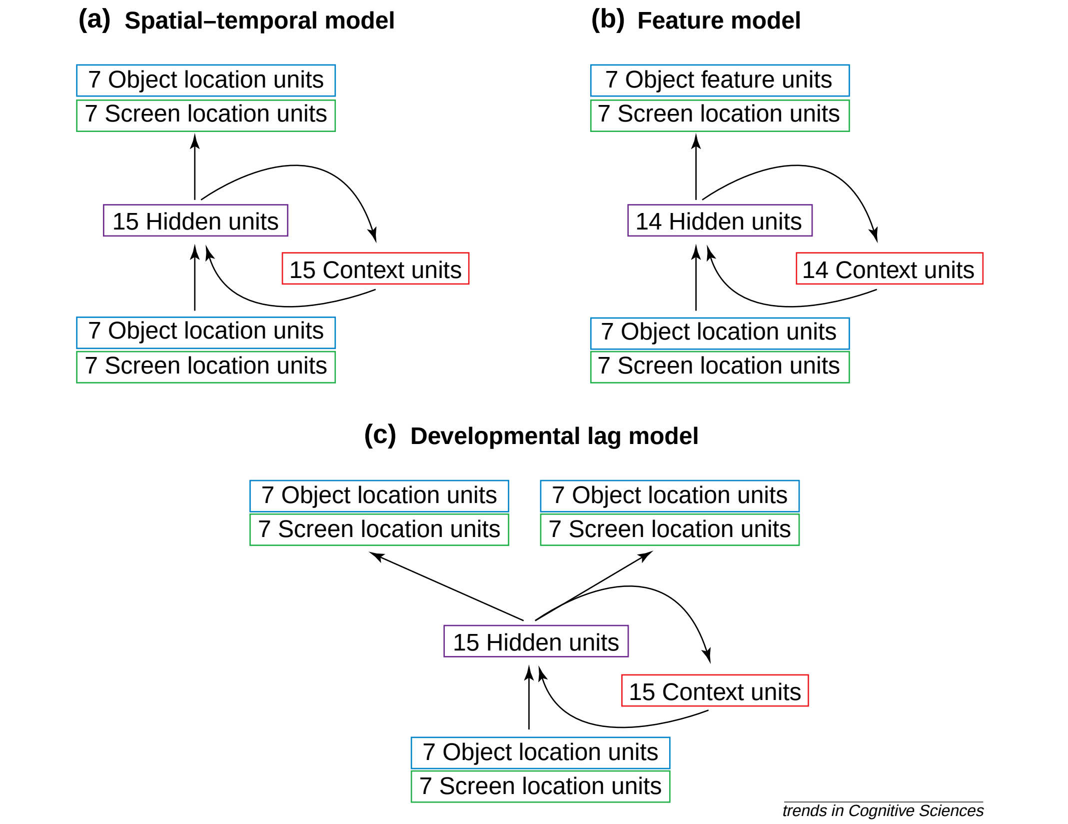

Fig. I. Schemas of three different models designed to predict infant object-directed behaviours. (a) The spatial–temporal model – predictive tracking of location, **(b)** the feature model – predictive tracking of features and **(c)** the developmental lag model – comparing predictive tracking with predictive reaching.

One of the first attempts to use PDP methods to model infant object-directed behaviours involved using partially recurrent autoencoder networks (Ref. a). In these models, connectionist networks received information about a simple environment, consisting of a stationary object and a moving screen, via a simple input retina. The networks were trained using backpropagation to predict the next perceptual input. Three models, corresponding to three different simulations, were reported (Fig. I).

In model 1 (Fig. Ia, the spatial–temporal model), the location of a ball and a moving occluding screen were presented to the network retina. The retina consisted of 14 units. Seven units were used to code each of the seven possible locations of a ball. The other seven units were used to code the shifting location of an occluding screen hypothesized to translate across the retina. As the screen moved across the position occupied by the ball, the input unit coding the ball location was turned off, to represent the fact that the ball was no longer visible. The network was trained to reproduce, on the output units, what the next input to the network would be. Knowledge (or at least internal representation) of hidden objects was evaluated by assessing the ability of the network to turn on the correct ball location unit when the ball was expected to reappear after occlusion, as compared to its tendency to activate that unit when no ball was expected to reappear. Only the unit coding the location of the stationary ball and the timeframe at which the ball was expected to reappear contributed to the evaluation of the model.

Model 2 (Fig. 1b, the feature model) was similar to model 1 except that feature information instead of location information was presented to the network. The prediction task was the same as in model 1, except that rather than responding with the predicted location of the ball when it was expected to reappear, the networks had to output the feature description of the object across the band of feature output units. A number of different objects could be represented, each by a unique feature representation distributed across the seven feature nodes.

Model 3 (Fig. Ic, the developmental lag model) was designed to model the developmental lag between infants' ability to predict the location of a hidden object and their ability to respond on that knowledge. This model adapted the spatial–temporal model by adding a second output route. The task in the two routes was identical (and identical to that in model 1). One route was intended to model the development of expectation abilities (as measured by looking time in infants) and the other route was intended to model the development of reaching abilities. To capture the lag between looking and reaching, the learning rate in the reaching route was reduced to one-tenth that in the predictive route and training on the reaching route was delayed until the prediction route had partially learned its task.

Reference

a Munakata, Y. *et al.* (1997) Rethinking infant knowledge: towards an adaptive process account of successes and failures in object permanence tasks. *Psychol. Rev.* 104, 686–713

representations that causes the developmental lag, this model suggests that surprise-based tasks requiring this integration to elicit a surprise response would also show a lag with respect to tasks that only required access to a single representation. The model shifts the account of the origins of the developmental lag from one involving differences in

Box 4. An object-processing model

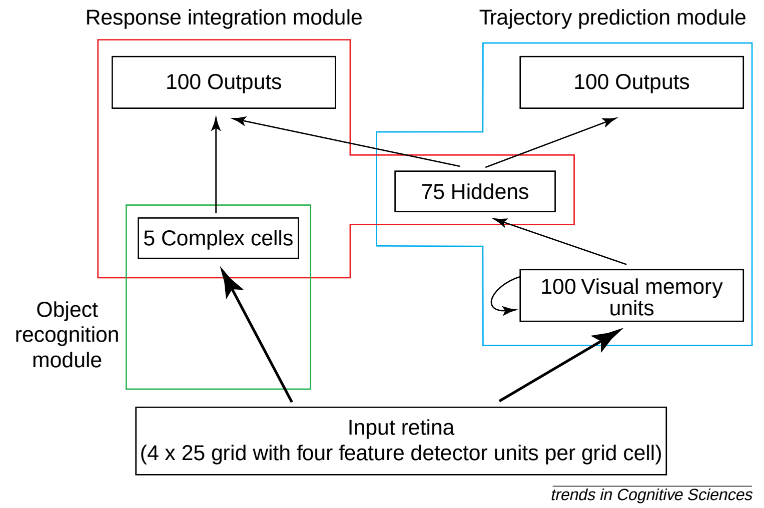

Fig. I. Schema of an object-processing model. (Adapted from Ref a.)

Figure I shows a schematic outline of an alternative connectionist model of the developmental lag in infant object-directed behaviours (Ref. a). This model has a modular architecture and is loosely based on the dual route visual processing hypothesis (Refs b,c). A coloured line encloses each functional module. Some units are shared by two modules (e.g. the 75 hidden units are shared by the response integration and trajectory prediction modules) and serve as a gateway for information between the modules.

The 'object recognition module' generates a spatially invariant representation of the object by using a modified version of the unsupervised learning algorithm developed by Foldiak (Refs d,e). This algorithm belongs to the family of competitive learning algorithms and exploits the fact that an object tends to be contiguous with itself at successive temporal intervals. Thus, two successive images will probably be derived from the same object. At the end of learning each complex cell becomes associated with a particular feature combination wherever it appears on the retina. This module is loosely based on ventral route processing.

The 'trajectory prediction module' uses a partially recurrent, feed-forward network trained with the backpropagation learning algorithm. This network learns to predict the next instantaneous, retinal position of the object. The internal representations it develops are determined both by the computational properties of the associative mechanisms in the network and by the spatial–temporal prediction task it is engaged in. This module is loosely based on dorsal route processing.

The output of the 'response integration network' corresponds to the infant's ability to co-ordinate and use the information it has about object position and object identity. This network integrates the internal representations generated by other modules (i.e. the feature representation at the complex cell level and spatial–temporal representation in the hidden unit layer) as required by a retrieval response task. It consists of a single layered perceptron whose task is to output the same next position as the prediction network for two of the objects and to inhibit any response (all units set to 0.0) for the other two objects. This reflects the fact that infants do not retrieve (e.g. reach for) all objects. In general, infants are not asked or rewarded for search. The experimental set-up relies on spontaneous search by the infant. Some objects are desired (e.g. *sweet*) whereas others are not desired (e.g. *sour*). Any voluntary retrieval response will necessarily require the processing of feature information (to identify the object as a desired one) as well as trajectory information (to localize the object). Surprise is modelled by a mismatch between the information stored in an internal representation and the new information arriving from the external world. The functions of this module can be loosely related to those of the frontal lobes.

Early mastery of surprise tasks that claim to show the co-ordination of position and feature information (Ref. f) have – on close scrutiny – provided evidence only for the use of positional information in conjunction with size or volume information. Both size and volume are spatial dimensions that are encoded by the dorsal route requiring access to only a single cortical route. Note that early surprise responses can arise from feature violations, from spatial–temporal violations and even from both types of violation arising concurrently and independently, but not from a violation involving the integration of feature and spatial–temporal information concerning an occluded object. The model predicts that infants will show a developmental lag not just on manual search tasks but also on surprise tasks that involve such integration. Conversely, the model suggests that infants will show early mastery of response tasks that do not require the integration of information across cortical representations.

The developmental lag for occluded objects arises as a natural consequence of the associative learning process. Internal object representations developed over the complex cells and the hidden units persist when the object passes behind the screen, but decay with time. Hence, activation levels drop when the object is occluded. The learning algorithm updates network weights in proportion to the activation level of the sending unit. For an identical error signal, the weight updates are smaller when the object is hidden given the lower activation of the sending units. Consequently, it will take longer to arrive at an equivalent level of learning for hidden than for visible objects. This outcome is not unique to the learning algorithm used in the current model. It will arise in any learning mechanism that updates weights in proportion to the sending unit activation, providing a clear example of how developmental behaviours are constrained by microlevel mechanisms.

References

a Mareschal, D. *et al.* (1999) A computational and neuropsychological account of object-oriented behaviors in infancy. *Dev. Sci.* 2, 306–317

b Milner, A.D. and Goodale, M.A. (1995) *The Visual Brain in Action*, Oxford University Press

c Van Essen, D.C. *et al.* (1992) Information processing in the primate visual system: an integrated systems perspective. *Science* 255, 419–423

d Foldiak, P. (199l) Learning invariance from transformation sequences. *Neural Comput.* 3, 194–200

e Foldiak, P. (1996) Learning constancies for object perception. In *Visual Constancies: Why Things Look As They Do* (Walsh, V. and Kulikowski, J., eds), pp. 144–172, Cambridge University Press

f Baillargeon, R. (1993) The object concept revisited: new directions in the investigation of infants' physical knowledge. In *Visual*

Box 5. How far can models go?

A major challenge to any learning account of infant object-directed behaviours (whether connectionist or otherwise) is the argument that no kind of empirical evidence or experience is sufficient to get the infant from a state in which the endurance of objects is not represented, to a state in which it is represented (Ref. a,b). There are two possible interpretations of this argument. One is related to the 'poverty of the stimulus argument' against radical empiricism. The other is related to Fodor's paradox of concept learning.

The 'poverty of the stimulus' argument was most effectively put forward by Chomsky in response to Skinner's account of language acquisition (Ref. c). It states that it is not possible for an unconstrained inductive learner to acquire a particular target grammar within reasonable time. This argument has been wielded against connectionist models of language and cognitive development in general (Ref. d). However, it is important to understand that the 'poverty of the stimulus' argument holds for all inductive learning systems and in all domains (Ref. e) (its application is not unique to connectionism, nor to learning about objects). This is why most contemporary scholars of learning (whether studying learning in children or machines) believe that the key to understanding cognitive development is to identify the nature of the constraints on the learner that will allow knowledge to emerge (Refs e,f).

Fodor's paradox claims that an inductive learner can never acquire any truly novel concept (Ref. g). Indeed, in order to test the domain of applicability of a concept, the inductive learner must be able to represent that concept prior to having identified it. Hence, any learning simply involves the recombination of existing representational tokens in a system. Although this may be true of inductive learning systems, it is not true of systems that increase their representational power in response to environmental pressures. Such systems include neural networks that construct their own architecture as part of learning and development (Ref. h). Such models distinguish between learning (weight adjustment in the network) and development (adding units to the network) and echo the cry for a similar distinction to be made when studying cognitive development (Ref. i).

References

- **a** Spelke, E.L. (1991) Physical knowledge in infancy: reflections on Piaget's theory. In *The Epigensis Mind. Studies in Biology and Cognition* (Carey, S. and Gelman, R., eds), pp. 154–198, Lawrence Erlbaum
- **b** Spelke, E.S. (1994) Initial knowledge: six suggestions. *Cognition* 51, 431–445
- **c** Chomsky, N.A. (1959) Review of verbal behavior. *Language* 35, 26–58
- **d** Marcus, G.F. (1998) Can connectionism save constructivism? *Cognition* 66, 153–182
- **e** Mitchell, T.M. (1997) *Machine Learning*, McGraw-Hill
- **f** Elman, J.L. *et al.* (1996) *Rethinking Innateness: Connectionism in a Developmental Framework*, MIT Press
- **g** Fodor, J. (1980) Fixation of belief and concept acquisition. In *Language and Learning: The Debate between Chomsky and Piaget* (Piatelli-Palmarini, M., ed.), pp. 143–149, Harvard University Press
- **h** Mareschal, D. and Shultz, T.R. (1996) Generative connectionist networks and constructivist cognitive development. *Cognit. Dev.* 11, 571–603
- **i** Liben, L.S. (1987) *Learning and Development: Conflict or Congruence?*, Lawrence Erlbaum

the response modality to one involving the co-ordination of multiple object representations. It predicts that developmental lags can occur even within the same response modality (e.g. looking time) and are not dependent on measuring responses across two different response modalities. Of course, this minimal lag could also be compounded by modality-specific developmental differences[12].

Future directions

These models are by no means complete. Much of infant behaviour still remains to be explained by any computational model – whether symbolic or connectionist. The existing connectionist models illustrate how object representations (that persist beyond direct perception) can emerge through interactions with the environment. Both sets of connectionist models provide an account of how a developmental lag can emerge between infant behaviour as tested on visual preference tasks and their behaviour when tested with manual retrieval tasks. Moreover, the final model[44] suggests how the need for the integration of cortically separable object representations might lead to infant behaviours in reaching and visual preference tasks considered separately.

However promising these models appear, both the connectionist and symbolic models still face a number of important challenges (Box 5). Not the least of these challenges is to see how the models (and the principles they embody) cope with being extended to the multitude of tasks in which infants show knowledge of hidden objects. For example, although the computational principles embodied by these models bear on the more recent infant findings reported above, none of the connectionist models directly models the task paradigms used to assess object individuation.

The original question posed by Piaget – of how a complete understanding of hidden objects can be acquired through interactions with the environment – remains unanswered. What the computational models contribute to this debate is a demonstration that internal representations of hidden objects that are perceptually independent can be acquired through learning. How these map onto the full and rich adult concept of object permanence is still unclear[56]. Ironically, Piaget's suggestion that object knowledge develops through the integration and co-ordination of multiple representations (perceptual–motor schema in his language) may not be so far off the mark.

Outstanding questions

- How does the infant's emerging competence with visible objects impact on his or her abilities to respond to hidden objects?
- Can a concept of 'object permanence' be acquired through learning? If so, then what kind of events are required for learning?
- In what way does the initial familiarization phase of many expectation-based studies bias infants' abilities to attend to certain object properties and not others?
- How can cortical development during infancy inform this debate (especially with regards to cortical object representations)?
- Does the distinction between 'voluntary' and 'involuntary' responses bear on the level of knowledge demonstrated by infants at different ages?

Acknowledgements

I am very grateful to Andrea Aguiar, Mark Johnson, Scott Johnson, and Teresa Wilcox for helpful comments on an earlier draft of this paper. This work was supported by a grant from the Sackler Institute.

References

1 Piaget, J. (1952) *The Origins of Intelligence in the Child*, International Universities Press

2 Piaget, J. (1954) *The Construction of Reality in the Child*, Basic Books

3 Harris, P. (1983) Infant cognition. In *Handbook of Child Psychology* (Vol. 2, 4th edn) (Mussen, P.H., ed.), pp. 689–781, John Wiley & Sons

4 Baillargeon, R. *et al*. (1985) Object permanence in 5-month-old infants *Cognition* 20, 191–208

5 Baillargeon, R. (1987) Object permanence in 3.5- and 4.5-month-old infants. *Dev. Psychol.* 23, 655–664

6 Baillargeon, R. (1993) The object concept revisited: new directions in the investigation of infants' physical knowledge. In *Visual Perception and Cognition in Infancy* (Granrud, C.E., ed.), pp. 265–315, Lawrence Erlbaum

7 Spelke, E.S. *et al.* (1992) Origins of knowledge. *Psychol. Rev.* 99, 605–632

8 Spelke, E.S. *et al.* (1994) Early knowledge of object motion: continuity and inertia. *Cognition* 51, 131–176

9 Spelke, E.S. (1994) Initial knowledge: six suggestions. *Cognition* 51, 431–445

10 Von Hofsten, C. (1989) Transition mechanisms in sensori-motor development. In *Transition Mechanisms in Child Development: The Longitudinal Perspective* (De Ribaupierre, A., ed.), pp. 223–259, Cambridge University Press

11 Fischer, K.W. and Bidell, T. (1991) Constraining nativist inferences about cognitive capacities. In *Epigenesis of Mind: Essays on Biology and Cognition* (Carey, S. and Gelman, R., eds), pp. 99–126, Lawrence Erlbaum

12 Munakata, Y. *et al*. (1997) Rethinking infant knowledge: towards an adaptive process account of successes and failures in object permanence tasks. *Psychol. Rev.* 104, 686–713

13 Diamond, A. (1993) Neuropsychological insights into the meaning of object concept development. In *Brain Development and Cognition* (Johnson, M.H., ed.), pp. 208–247, Blackwell

14 Diamond, A. (1985) The development of the ability to use recall to guide action, as indicated by infants' performance on A-not-B. *Child Dev.* 56, 868–883

15 Thelen, E. and Smith, L.B. (1994) *A Dynamic Systems Approach to the Development of Cognition and Action*, MIT Press

16 Baillargeon, R. *et al*. (1989) Location memory in 8-month-old infants in a non-search AB task: further evidence. *Cognit. Dev.* 4, 345–367

17 Baillargeon, R. (1998) Infants' understanding of the physical world. In *Advances in Psychological Science* (Vol. 1) (Sabourin, M. *et al*., eds), pp. 305–371, Psychology Press

18 Aguiar, A. and Baillargeon, R. (1999) 2.5-month-olds' reasoning about when objects should and should not be occluded. *Cognit. Psychol.* 39, 116–157

19 Baillargeon, R. (1999) Young infants' expectations about hidden objects: a reply to three challenges. *Dev. Sci.* 2, 115–163

20 Wilcox, T. and Baillargeon, R. (1998) Object individuation in infancy: the use of featural information in reasoning about occlusion events. *Cognit. Psychol.* 17, 97–155

21 Artberry, M.E. (1997) Perception of object properties over time. In *Advances in Infancy Research* (Vol. 11) (Rovee-Collier, C. and Lipsitt, L.P., eds), pp. 219–268, Ablex

22 Baillargeon, R. and DeVos, J. (1991) Object permanence in young infants: further evidence. *Child Dev.* 62, 1227–1246

23 Spelke, E.S. *et al.* (1995) Spatiotemporal continuity, smoothness of motion and object identity in infancy. *Br. J. Dev. Psychol.* 13, 1–30

24 Wilcox, T. *et al*. (1996) Location memory in healthy preterm and full-term infants. *Infant Behav. Dev.* 189, 309–323

25 Wynn, K. (1992) Addition and subtraction in human infants. *Nature* 358, 749–750

26 Wynn, K. (1995) Infants possess a system of numerical knowledge. *Curr. Dir. Psychol. Sci.* 4, 172–177

27 Simon, T.J. *et al*. (1995) Do infants understand simple arithemtic? A replication of Wynn. *Cognit. Dev.* 10, 253–269

28 Xu, F. and Carey, S. (1996) Infants' metaphysics: the case of numerical identity. *Cognit. Psychol.* 30, 111–153

29 Leslie, A.M. *et al.* (1998) Indexing and the object concept: developing 'what and 'where' systems. *Trends Cognit. Sci.* 2, 10–18

30 Wilcox, T. and Baillargeon, R. (1998) Object individuation in young infants: further evidence with an event-monitoring task. *Dev. Sci.* 1, 127–142

31 Wilcox, T. (1999) Object individuation: infants' use of shape, size, pattern, and color. *Cognition* 72, 125–166

32 Needham, A. and Baillargeon, R. (2000) Infants' use of featural and experiential information in segregating and individuating objects: a reply to Xu, Carey and Welch (1999). *Cognition* 74, 255–284

33 Xu, F. and Carey, S. (2000). The emergence of kind concepts: a rejoinder to Needham and Baillargeon (2000). *Cognition* 74, 285–301

34 Johnson, S.P. (1998) Object perception and object knowledge in young infants: a view from studies of visual development. In *Perceptual Development* (Slater, A., ed.), pp. 211– 245, Psychology Press

35 Haith, M.M. and Benson, J.B. (1998) Infant cognition. In *Handbook of Child Psychology* (Vol. 2, 5th edn) (Damon, W., series ed.; Kuhn D. and Siegler R., vol. eds), pp. 199–254, John Wiley & Sons

36 Haith, M.M. (1999) Some thoughts about claims for innate knowledge and infant physical reasoning. *Dev. Sci.* 2, 153–156

37 Haith, M.M. *et al*. (1993) The formation of expectation in early infancy. In *Advance in Infancy Research* (Rovee-Collier, C. and Lipsitt, L.P., eds), pp. 53–76, Ablex

38 Bogartz, R.S. *et al*. (1997) Interpreting infant looking: the event set × event set design. *Dev. Psychol.* 33, 408–422

39 Rivera, S.M. *et al.* (1999) The drawbridge phenomenon: representational reasoning or perceptual preference *Dev. Psychol.* 35, 427–435

40 Haith, M.M. *et al.* (1994) *The Development of Future-Oriented Processes*, University of Chicago Press

41 Haith, M.M. (1998) Who put the cog in infant cognition? Is rich interpretation too costly? *Infant Behav. Dev.* 21, 167–179

42 Spelke, E.S. (1998) Nativism, empiricism, and the origins of knowledge. *Infant Behav. Dev.* 21, 181–200

43 Berthenthal, B.I. (1996) Origins and early development of perception, action, and representation. *Annu. Rev. Psychol*. 47, 431–459

44 Mareschal, D. *et al*. (1999) A computational and neuro-psychological account of object-oriented behaviors in infancy. *Dev. Sci.* 2, 306–317

45 Smith, L.B. (1999) Do infants possess innate knowledge structures? The con side. *Dev. Sci.* 2, 133–144

46 Rutkowska, J.C. (1993) *The Computational Infant: Looking for Developmental Cognitive Science*, Harvester Wheatsheaf

47 Simon, T.J. and Halford, G.S. (1995) *Developing Cognitive Competence: New Approaches to Process Modeling,* Lawrence Erlbaum

48 Prazdny, S. (1980) A computational study of a period of infant object–concept development. *Perception* 9, 125–150

49 Luger, G.F. *et al*. (1983) A model of the development of the early object concept. *Perception* 12, 21–34

50 Luger, G.F. *et al*. (1984) Modeling the stages of the identity theory of object concept development in infancy. *Perception* 13, 97–115

51 Drescher, G.L. (1991) *Made-up Minds. A Constructivist Approach to Artificial Intelligence,* MIT Press

52 Fischer, K.W. and Bidell, T. (1991) Constraining nativist inferences about cognitive capacities. In *Epigenesis of Mind: Essays on Biology and Cognition* (Carey, S. and Gelman, R., eds), pp. 99–126, Lawrence Erlbaum

53 Milner, A.D. and Goodale, M.A. (1995) *The Visual Brain in Action*, Oxford University Press

54 Milner, A.D. (1998) Streams and consciousness: visual awareness and the brain. *Trends Cognit. Sci.* 2, 25–30

55 Passingham, R.E. (1993) *The Frontal Lobes and Voluntary Action*, Oxford University Press

56 Scholl, B.J. and Leslie, A.M. (1999) Explaining the infant's object concept: beyond the perception/cognition dichotomy. In *What is Cognitive Science?* (Lepore, E. and Pylyshyn, Z., eds), pp 26–73, Blackwell