Opinion TRENDS in Cognitive Sciences Vol.9 No.3 March 2005

Methodological challenges for understanding cognitive development in infants

Richard N. Aslin and József Fiser

Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY 14627, USA

Studies of cognitive development in human infants have relied almost entirely on descriptive data at the behavioral level – the age at which a particular ability emerges. The underlying mechanisms of cognitive development remain largely unknown, despite attempts to correlate behavioral states with brain states. We argue that research on cognitive development must focus on theories of learning, and that these theories must reveal both the computational principles and the set of constraints that underlie developmental change. We discuss four specific issues in infant learning that gain renewed importance in light of this opinion.

Introduction

Over the past 40 years, research on human cognitive development has revealed a set of surprisingly sophisticated competencies in infancy, given the seemingly immature and incompetent appearance of the newborn [1]. These findings emerged as a result of increasingly subtle and sensitive methods, most of which rely on behavioral responses such as visual fixations and eye movements, sucking, and a variety of psychophysiological recordings (e.g. heart rate) [2]. In the past decade, there has been a resurgence of interest in obtaining more direct recordings of brain activity (sensory-evoked potentials, event-related potentials, EEG spectra) and their hemodynamic correlates (PET, fMRI, near-infrared spectroscopy) [3,4]. These new cognitive neuroscience methods have brought infancy research much closer to the neurobiological bases of behavioral development and reduced the gap in methodologies and viewpoints between traditional sensorimotor neuroscience and developmental science.

Despite these advances, we believe that several key research strategies must be followed for substantial progress to continue. First, experimental studies of cognitive development in infants must shift from predominantly single-parameter tests that assess the presence/absence of a capacity at a particular age to multiple-parameter tests that assess the full range of an underlying mechanism across age. Second, the role of learning in accounting for cognitive development must be expanded and a computational learning framework must be developed to bring studies of infant perceptual and cognitive processing to the level of sophistication emerging in research on mature sensory and motor behavior [5,6]. Third, as learning mechanisms must be constrained to enable them to plausibly acquire the type of information that complex perceptual and cognitive processes require, theories of learning must be directed to studies that reveal these constraints. Fourth, rather than demonstrating the localization of neural activity in the nervous system, the new methods of developmental cognitive neuroscience should focus on testing directly the computational framework for how information is learned in the brain.

Here, we review the relevant data acquisition paradigms and the computational requirements for developing an adequate framework of infant learning, and we discuss how such a framework could be further constrained by the methods of cognitive neuroscience.

The power and limitations of behavioral data

Presence–absence paradigms

Nearly all studies of cognitive and language development in human infants use paradigms that can only provide group data. This is in part because of the limited duration of infant cooperation (10–15 min) and in part because infant performance is more susceptible to fluctuations in attention and motivation than adults. As a result, each infant typically provides a very few (1–4) data points per condition and statistical comparisons are made between two treatment groups (or an experimental and a control group). The paradigms used with infants (see Box 1) essentially provide a Yes–No answer to a given research question, such as whether ability X is present at a given age. A limitation of Yes–No data is that, even with large sample sizes, it is difficult to determine which, among the dozens of potentially relevant stimulus parameters in a given experiment, enabled the infants (as a group) to exhibit positive evidence of ability X. To clarify the roles of these parameters, additional experiments must be conducted with this binary testing method to obtain a coherent set of outcomes.

In addition, many behavioral paradigms (Box 1) are not only limited to providing a Yes–No answer to a given research question but also the difference (e.g. from a control group) can be in either of two directions (e.g. familiarity or novelty) [7–10]. This bidirectional outcome is not, in principle, a problem if the desired conclusion is that infants can discriminate between two

Corresponding author: Aslin, R.N. (aslin@cvs.rochester.edu).

Available online 29 January 2005

www.sciencedirect.com 1364-6613/$ - see front matter doi:10.1016/j.tics.2005.01.003

Box 1. Behavioral paradigms and their measurement outcomes

Measures of infant cognition that rely on looking times to visual or auditory stimuli fall into three categories: preference, familiarization/habituation and violation of expectancy. Sample paradigms are shown in Figure I. In each paradigm, infants are exposed to a very small number of test trials, in some cases after being pre-exposed to several stimuli and in other cases with no pre-exposure trials (see Table I). Trials are either of fixed duration or of variable duration controlled by the infant (e.g. fixations directed away from the stimulus for less than two consecutive seconds do not result in termination of a trial). For visual studies, looking times to two or more visual stimuli serve as the dependent variable. For auditory studies, looking times to a uniform visual stimulus presented in the presence of two or more auditory stimuli serve as the dependent variable.

Table I. The limited number of data points collected from each infant

Paradigm	Number of pre-test trials	Number of test trials
Preference	0	2–12
Familiarization/Habituation	6–20	2–6
Violation of expectancy	0–4	2–6

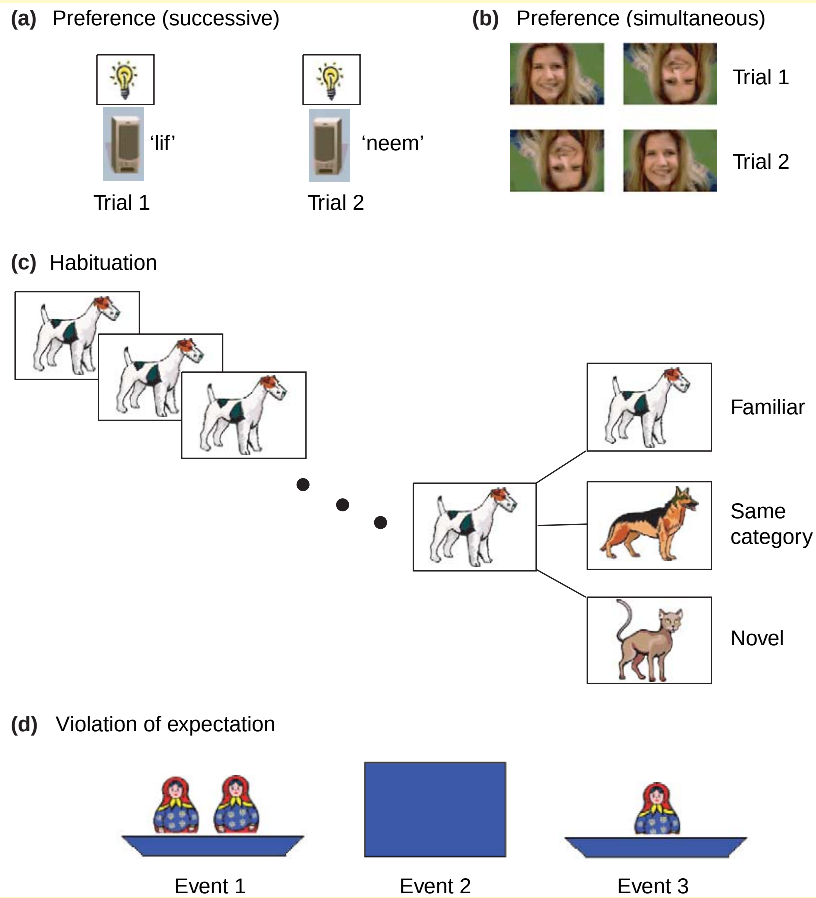

Figure I. The three major paradigms used to measure cognitive development in human infants. Preference paradigms use either successive **(a)** or simultaneous **(b)** test trials, in which a visual stimulus serves to elicit fixations. The duration of looking to the visual stimulus in the presence of two different auditory stimuli **(a)**, or to the two different visual stimuli themselves **(b)**, serves as the measure of preference. If looking times to the two types of test trials are different from each other, discrimination as well as preference can be inferred. Habituation paradigms **(c)** repeat a single visual stimulus (or a visual–auditory compound) until looking time declines to a preset criterion. Alternatively, familiarization can be of a fixed duration. Two (or more) types of test trials are then presented to assess differential recovery of looking time. Significant recovery from the last habituation trial to a test trial (compared to a no-change control) indicates discrimination. Greater recovery to one of two types of test trials reflects a familiarity or novelty preference with respect to the habituation stimulus. In a Multiple Habituation variant of this paradigm, two or more different stimuli are presented during the habituation phase, followed by test trials that are the same or different from the habituation stimuli along one common dimension. If there is no post-habituation recovery to novel test stimuli that share the common dimension, then infants are inferred to have categorized along that dimension. 'Violation of expectancy' paradigms present a series of visual events that are judged by adults to be 'impossible' because they contain properties that violate some physical or cognitive principle. For example, in **(d)** two objects are hidden behind an occluder and then a test trial reveals a single object when the occluder is removed. Longer looking times to impossible over possible test events provide evidence of some level of underlying 'knowledge'.

stimuli. But it becomes an interpretive problem when the direction of the effect is crucial to drawing a particular conclusion and cannot be predicted in advance. This can lead to post hoc explanations in a given experiment to account for a positive effect. Moreover, when a negative effect is obtained, it is not clear if it is a true failure (i.e. the absence of ability X) or rather a false negative resulting from a 'balance' of bidirectional outcomes when summed across infants in grouped data. As the novelty versus familiarity distinction is a crucial one in infant learning, this uncertainty becomes a major obstacle to gaining precise information using the paradigms of Box 1.

Stimulus–response (S–R) monotonicity

In contrast to the above 'cognitive' paradigms, there is a longstanding tradition in psychophysics and motor control that relies on the collection of many data points from each subject under three or more stimulus conditions (multiple S–R paradigms), and these paradigms have been transferred successfully to infant research [11,12]. For example, visual acuity can be assessed in individual infants by measuring fixations to stripes that vary in width, thereby generating a psychometric function relating a stimulus variable (stripe-width) to a response (percent detection). Moreover, because every test trial can be treated as independent (except for generalized fatigue across trials), a given infant can provide dozens of trials either within or across sessions. Although group data are often the preferred way of reporting infant performance, even when sufficient within-subject data are available, having access to within-subject data makes it possible to avoid the inappropriate collapsing of data across subjects, which can erroneously characterize changes in performance with age.

Yes–No paradigms are less powerful than multiple S–R paradigms in drawing conclusions about underlying mechanisms because there is no *function* relating the putative stimulus variable to the measured response. For example, we cannot conclude in a violation-of-expectancy paradigm that a doubling of a novelty preference to condition A versus condition B means that A is twice as discriminable as B. Rather, Yes–No paradigms use quantitative data (e.g. looking times) to draw qualitative conclusions. In addition, Yes–No paradigms rely on the appropriate control of all irrelevant stimulus variables, rather than observing the systematic outcome of altering the putative stimulus variable through three or more levels.

The dilemma facing researchers who want to assess cognitive and language development in infants is that detection of a stimulus (as in a visual acuity study) is not typically the goal of their research questions. Instead, they seek to characterize what stimulus information *can* be used by infants in a specific task or how stimulus information is *learned*. This in turn means that infants must be exposed to a 'corpus' of stimulus information (either before or during the experimental session) and tested for their ability to extract from that corpus one or more types of information. In our judgment, studies of learning require more than the kinds of qualitative data provided by Yes–No paradigms because a proper understanding of underlying mechanisms demands quantitative results to test competing models. This is not to say that Yes–No paradigms are always inconclusive, but they tend to be very inefficient (requiring large sample sizes) and subject to false-negative outcomes.

In summary, future progress in studying infant cognitive development will benefit from paradigms that are more akin to those used with adults. Rich datasets from within subjects are needed so that the range of performance affected by variation along a stimulus dimension of interest can be revealed. Current Yes–No paradigms are both inefficient and potentially misleading as a characterization of development. However, even if rich datasets were available from infants, the field needs to confront the mechanisms of learning to understand how information is acquired and represented in the brain.

Information extraction: system-level mechanisms of learning

Unsupervised but not unconstrained

We propose that the fundamental limitation on progress in research on infant cognitive development is at the level of a computational theory of infant learning. By infant learning we refer to the full spectrum of knowledge acquisition: from low-level perceptual processes, such as understanding the three-dimensional layout of a scene, to high-level cognitive processes, such as language acquisition. Computational studies of infant learning must confront the following dilemma. Learning mechanisms can be made more efficient by the availability of specific feedback to direct the information–extraction process. This suggests that for domains of cognitive development where acquisition occurs rapidly, the most plausible framework for infant learning is that of supervised learning. As a result, researchers studying higher-level cognitive and language development initially used classical supervised methods to model infant development [13–15]. However, there is little evidence that specific feedback is available in the infant's natural environment to guide lower-level sensory and motor learning (other than crude after-the-fact feedback, as in reaching, that has no direct effect on the subcomponents of the reach). This suggests that either reinforcement learning [16], which relies on indirect feedback, or unsupervised learning [17,18], which relies on information structure rather than feedback *per se*, is a more plausible mechanism for infant learning. Thus, infant learning can be characterized by different 'modes': supervised learning, reinforcement learning, unsupervised learning and, in rare cases, one-shot learning. We suggest that a successful learning framework for infant development must encompass all these modes of learning because lower-level sensory and motor learning is intimately integrated into the development of higher-level cognitive functions. Without doubt, the underlying neural mechanisms by which these modes of learning operate and the biophysical basis of learning and plasticity in the brain (e.g. long-term potentiation, LTP, and long-term depression, LTD) may be diverse. However, we believe that a viable goal is to search for a unified system-level explanation of how information is extracted and represented in the brain that can account for phenomena at both lower and higher levels of cognition as well as across different domains.

A second fundamental problem with attempting to develop a framework of learning that can handle large complex problems, such as visual perception or language, is that a powerful mechanism for acquiring information must be able to deal with the exponentially large number of potentially relevant correlations available in the input. For example, in randomly ordered strings of four different elements, there are 12 possible unique pairs. But when the strings are doubled to eight different elements, the number of possible pairs increases more than fourfold to 56. This exponential increase requires an equally expanding number of sensory examples to learn the relevant pairs that are meaningful for the organism. However, under natural situations there is no realistic opportunity for infants to experience so many examples. This is called the curse of dimensionality [19] or the computational explosion problem [20], and it is particularly insidious when considering not just adjacent pairs of elements but also nonadjacent pairs and higher-order combinations of more than two elements in any sensory or cognitive domain.

Despite the presence of the computational explosion problem, human learners have been shown to be sensitive to the statistical probabilities that define the underlying structures embedded in both sequential strings (visual and auditory) and spatial arrangements of visual elements [18,21,22]. Moreover, they automatically extract a variety of statistics, including conditional probabilities [23,24], which provide a superior metric of predictability than simple counts of the relative frequency of element co-occurrences. Importantly, this statistical learning mechanism is robust and rapid, and operates in an unsupervised mode in both adults and infants [25–27]. In addition, there is rapidly accumulating evidence that in human adults low-level perception and action is optimal in a Bayesian probability sense [5,6]. This raises the possibility that infants' internal representations of visual, auditory and other sensory inputs develop based on a general statistical learning method.

Computational bootstrapping: less is more

How can we reconcile the computational explosion problem with the empirical evidence that statistical learning operates rapidly and efficiently in both adults and infants? From adult experiments, the solution seems to involve a mechanism that generates an internal representation that only partially codes the sensory input while it reduces the redundancy of the code [28,29]. That is, when correlations among groups of elements in the input contain embedded correlations among smaller subgroups of elements, the redundant subgroups are eliminated (or downweighted) in the corresponding feature representation generated in the brain. This resembles the behavior of standard unsupervised learning algorithms that uncover the underlying causes of complex modular inputs [30,31]. However, the only way to use such a method on large real-world problems is to add a set of constraints to implicitly guide learning to acquire a minimally sufficient rather than a complete representation of the input. For infants, these constraints could be endogenous, by building in specialized mechanisms, or exogenous, by having caregivers reduce the complexity of the input itself or otherwise provide their infant with an environment that has a limited range of structures. There are examples of each of these types of constraints in different domains. For example, in language development mothers hyperarticulate their speech to infants [32] and embed words in short, grammatically simple, and prosodically salient utterances [33]. In vision, even newborns have a preference to attend to roughly face-like or slowly moving patterns [34,35].

An alternative to endogenous constraints that arise from specialized mechanisms are constraints that are the byproduct of general developmental immaturities. Originally proposed in the domain of language development, with the phrase 'less-is-more' [36], these constraints are based on the limited sensory and working-memory capacities of infants (or any naïve learner confronted with initially unfamiliar stimuli [37]; see Box 2). If sensory immaturities and limited working memory reduce the number of elements that can potentially be related via a statistical learning mechanism, then the computational explosion problem can be reduced or eliminated. Once 'simple' element correlations are extracted, they can in turn serve as the elements for building higher-order features. In other words, developmental immaturities allow infants to bootstrap their way to a feature hierarchy that avoids the curse of dimensionality.

An important outcome of this computational bootstrapping process is that, if the representational power of the initial set of features is sufficiently rich, the developmental feature hierarchy enables infants (and adults) to recognize both familiar and novel events. Familiar events are recognized by noting the higher-order features, whereas novel events are recognized by noting the lower-order features (and the fact that they are organized in novel ways). For example, the perception of faces is a highly efficient process of extracting familiar lower-order features and organizing them into higher-order spatial configurations [41]. When the lower-order features are presented in a novel configuration (e.g. inverted), the higher-order recognition process is disrupted, but recognition is still possible if based on the lower-order features alone. If such bootstrapping is indeed the core mechanism for developing neural representations at both higher and lower levels, then a main objective of research on learning should be to identify the set of constraints that cooperate with the statistical learning process on each level to achieve this computational bootstrapping process.

Of course, the particular mechanisms by which information is learned and represented in the brain cannot be revealed entirely by data at the behavioral level. Behavioral data provide a rich source of information about *what* infants have learned, the *contexts* in which they learn, and the *constraints* that enable learning to occur both rapidly and efficiently. However, a direct neural assessment of how information is represented in the brain requires the methods of cognitive neuroscience.

The promise and prospects of cognitive neuroscience for development

Unaddressed by the foregoing behavioral evidence of statistical learning are the neural mechanisms that underlie it and the manner in which the output of statistical learning is represented in the brain. The same

Box 2. Limitations on working memory

A number of recent studies suggest that, despite impressive long-term memory, infants have remarkably poor short-term or working memory. When the duration of retention is a few seconds, only 3–4 visual objects can be recognized [38]. When visual objects must be updated in working memory every second, then capacity is 1–2 objects [39]. And when two objects are successively hidden behind an occluder for a few seconds, the capacity of short-term memory is only one [40].

can be said about every other cognitive and linguistic skill. At issue is how the techniques of cognitive neuroscience can help to resolve these questions of neural mechanism and representation.

The seduction of neurophrenology

In the absence of direct evidence for the neural mechanisms that enable a particular cognitive or learning process, which can only come from studies of animals, the primary goal of developmental cognitive neuroscience has been to determine *where* in the brain activity is present during a particular task. Even if we ignore for the moment the enormous technical problems associated with localizing regions of activation in the brain (see Box 3), it is not always clear what is gained by obtaining evidence of where neural activity is located. It is important to emphasize that neural activity that is not under experimental control is merely a correlate of behavior [52]. Thus, knowing that a particular region of the brain is active during a particular behavioral/cognitive state does not 'explain' that state. Without a theory for how the brain is organized and how information is represented in it, neural correlates of behavior are no more important than a second measure of behavior. In this respect, cognitive neuroscience shares many of the methodological problems of infancy research, where knowing when a particular ability emerges does not explain how the ability came about.

There are circumstances, however, under which researchers could benefit from knowing which regions of the brain are active. First, if the same region of the brain is active across a wide age range during a given task, then it is likely that the relationship between the task and the underlying neural architecture is maintained over this developmental period. Second, if a different region of the brain is active as information is acquired, and behavioral performance changes with continued exposure, then it is likely that the underlying neural representation has also changed. Unfortunately, these goals are difficult to attain when the data are in the form of Yes–No answers. For example, if there is a left-hemisphere (LH) advantage for linguistic materials in infancy as there is in adulthood, it is seductive to conclude that the LH specialization is developmentally invariant. But without a quantitative measure of the *degree* of LH advantage, it is possible that a modest bias in infancy develops into a strong bias in

Box 3. Recent methods for assessing infant brain function

The methods of cognitive neuroscience used with adults have recently been extended successfully to infants (see Figure II). Building on a longstanding subfield of EEG and event-related potential methods [42], ERP source localization [43] and EEG spectra [44,45] have shown great promise for localizing brain activation in infants during stimulus events. In addition, fMRI has been extended to the assessment of language areas of the infant brain [46], and near-infrared spectroscopy (NIRS), a through-the-skull optical imaging technique, has been used to assess visual [47], cognitive [48–50] and language [51] development in infants.

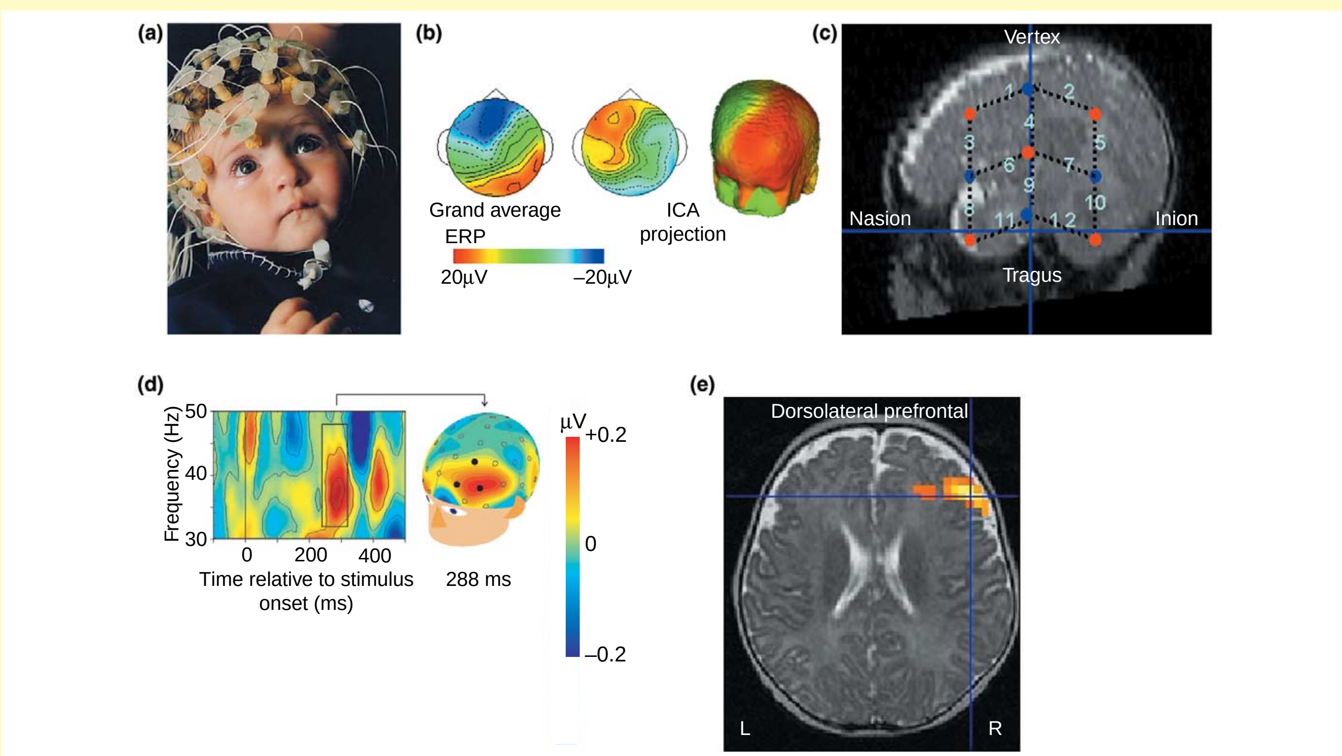

Figure II. Recent examples of the methods of cognitive neuroscience applied to human infants. **(a)** EGI 64-channel electrode net. (Photo courtesy of Charles Nelson.) **(b)** Independent components analysis (ICA) of ERPs to novel vs. familiar visual stimuli and the estimated cortical source localization. (Reproduced with permission from Ref. [43].) **(c)** Placement of near-infrared optical imaging probes used to detect hemodynamic responses from the left hemisphere of newborns presented with speech stimuli. (Reproduced with permission from Ref. [51].) **(d)** EEG spectral responses in the 40-Hz gamma band showing frontal activity to subjective contours. (Reproduced with permission from Ref. [44].) **(e)** BOLD fMRI responses from the frontal cortex to forward vs. backward speech stimuli. (Reproduced with permission from Ref. [46]).

adulthood, thereby altering the interpretation of the development of LH specialization.

The search for monotonicity in brain–behavior relationships

Our earlier point about the limitations of Yes–No paradigms in behavioral studies takes on added significance in studies of developmental cognitive neuroscience and highlights the common challenges faced by infancy researchers as they use any of the currently available dependent measures. Multiple S–R paradigms, like those used to study sensory and motor systems, provide a function relating a stimulus variable to a measured response. We believe that similar paradigms are essential in studies of brain–behavior relationships [53]. Consider the example of a specialized region in temporal cortex that is preferentially activated by faces and face-like objects. The so-called fusiform face area (FFA) may be an innate 'module' or it may emerge during development because of its capacity to represent objects that are highly familiar [54]. But without quantitative evidence for face and non-face discriminability and categorization [55], it will be very difficult to determine which of these competing theories of brain specialization is correct.

Conclusion

In this article, we have argued for four related points. First, experimental studies of cognitive development in infants require multiple S–R paradigms that help to clarify such basic issues as how infant learning is driven by a representational hierarchy that builds on familiar information to encode novel information. Second, a general learning framework of infant cognitive development is needed and it must be able to explain the unsupervised and supervised nature of learning, as well as the development of both low and high levels of information representation in the brain. Third, a promising framework includes unsupervised mechanisms that can bootstrap hierarchical representations, but such mechanisms require several evolutionarily built-in constraints on learning that enable complex tasks such as visual object recognition or language acquisition to be computationally tractable. Fourth, without a computational framework for how information is represented and learned in the brain, the localization of neural activity provided by the methods of cognitive neuroscience are of limited benefit. The goal of future studies in the field of cognitive development (see Box 4) should be to gather quantitative data using both behavioral and neural techniques so that detailed models of the mechanisms underlying learning can be evaluated.

Box 4. Questions for future research

- What determines when an infant pays more attention to novel versus familiar inputs? Is this dependence static or does it change as the input is repeated? Does it change with age?
- Is there a limit on the amount of information infants can take in at a glance? Is there a limit on the type of information infants can take in at a particular age? Can these limits be changed by suitable pre-training?
- How rich and detailed is the infant's representation of the environment? If the young infant's representation is impoverished, what allows it to improve with age? What does it mean that an infant 'cannot understand' an image?
- Is there a preset division of higher cortical areas or do the maps of functionally different areas develop and reorganize on demand during cognitive development? What principle determines which features will be represented in a more localized manner in the brain?
- Is there a common computational mechanism for the different types of learning, including fast one-shot learning, slow statistical learning, and the transfer of learning from one set of stimuli to another? How widespread is the mechanism (in terms of brain structures involved) by which new feature representations are acquired?

References

1 Kellman, P.J. and Arterberry, M.E. (1998) *The Cradle of Knowledge: Development of Perception in Infancy*, MIT Press
2 Gottlieb, G. and Krasnegor, N., eds (1985) *Measurement of Audition and Vision During the First Year of Postnatal Life: A Methodological Overview*, Ablex
3 Nelson, C.A. and Luciana, M. (2001) *Handbook of Developmental Cognitive Neuroscience*, MIT Press
4 deHaan, M. and Johnson, M.H. (2003) *The Cognitive Neuroscience of Development*, Psychology Press
5 Kersten, D. *et al*. (2004) Object perception as Bayesian inference. *Annu. Rev. Psychol.* 55, 271–304
6 Wolpert, D.M. and Ghahramani, Z. (2000) Computational principles of movement neuroscience. *Nat. Neurosci.* 3, 1212–1217
7 Rose, S.A. *et al*. (1982) Familiarity and novelty preferences in infant recognition memory: Implications for information processing. *Dev. Psychol.* 18, 704–713
8 Wagner, S.H. and Sakovits, L.J. (1986) A process analysis of infant visual and cross-modal recognition memory: Implications for an amodal code. In *Advances in Infancy Research* (Vol. 4) (Lipsitt, L. and Rovee-Collier, C., eds), pp. 195–217, Ablex
9 Hunter, M.A. and Ames, E.W. (1988) A multifactor model of infant preferences for novel and familiar stimuli. In *Advances in Child Development and Behavior* (Vol. 4) (Lipsitt, L.P., ed.), pp. 69–95, Academic Press
10 Roder, B.J. *et al*. (2000) Infants' preferences for familiarity and novelty during the course of visual processing. *Infancy* 1, 491–507
11 Teller, D.Y. (1979) The forced-choice preferential looking procedure: A psychophysical technique for use with human infants. *Infant Behav. Dev.* 2, 135–153
12 Dannemiller, J.L. (1998) A competition model of exogenous orienting in 3.5-month-old infants. *J. Exp. Child Psychol.* 68, 169–201
13 Elman, J.K. *et al*. (1996) *Rethinking Innateness: A Connectionist Perspective on Development*, MIT Press
14 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
15 Quinn, P.C. and Johnson, M.H. (1997) The emergence of perceptual category representations in young infants: A connectionist analysis. *J. Exp. Child Psychol.* 66, 236–263
16 Sutton, R. and Barto, A. (1998) *Reinforcement Learning: An Introduction*, MIT Press
17 Saffran, J.R. *et al*. (1996) Word segmentation: The role of distributional cues. *J. Mem. Lang.* 35, 606–621
18 Maye, J. *et al*. (2002) Infant sensitivity to distributional information can affect phonetic sensitivity. *Cognition* 82, B101–B111
19 Bellman, R. (1961) *Adaptive Control Processes: A Guided Tour*, Princeton University Press
20 von der Malsburg, C. (1995) Binding in models of perception and brain function. *Curr. Opin. Neurobiol.* 5, 520–526
21 Fiser, J. and Aslin, R.N. (2001) Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychol. Sci.* 12, 499–504
22 Fiser, J. and Aslin, R.N. (2002) Statistical learning of higher order temporal structure from visual shape-sequences. *J. Exp. Psychol. Learn. Mem. Cogn.* 28, 458–467
23 Aslin, R.N. *et al*. (1998) Computation of conditional probability statistics by 8-month-old infants. *Psychol. Sci.* 9, 321–324

24 Fiser, J. and Aslin, R.N. (2002) Statistical learning of new visual feature combinations by infants. *Proc. Natl. Acad. Sci. U. S. A.* 99, 15822–15826
25 Saffran, J.R. *et al*. (1996) Statistical learning by 8-month-old infants. *Science* 274, 1926–1928
26 Saffran, J.R. *et al*. (1999) Statistical learning of tone sequences by human infants and adults. *Cognition* 70, 27–52
27 Kirkham, N.Z. *et al*. (2002) Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition* 83, B35–B42
28 Fiser, J. and Aslin, R.N. (2002) Extraction of parts and wholes from multi-element scenes. *J. Vis.* 2, 558a
29 Fiser, J. and Aslin, R.N. (2003) Element predictability not high occurrence frequency determines feature learning from multi-element scenes. *J. Vis.* 3, 163a
30 Hinton, G.E. *et al*. (1995) The 'wake–sleep' algorithm for unsupervised neural networks. *Science* 268, 1158–1161
31 Dayan, P. *et al*. (1995) The Helmholtz machine. *Neural Comput.* 7, 889–904
32 Kuhl, P.K. *et al*. (1997) Cross-language analysis of phonetic units in language addressed to infants. *Science* 277, 684–686
33 Fernald, A. and Simon, T. (1984) Expanded intonation contours in mothers' speech to newborns. *Dev. Psychol.* 27, 209–221
34 Johnson, M.H. *et al*. (1991) Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition* 40, 1–19
35 Turati, C. (2002) Newborns preference for faces: What is crucial? *Dev. Psychol.* 38, 875–882
36 Newport, E.L. (1990) Maturational constraints on language learning. *Cogn. Sci.* 14, 11–28
37 Kersten, A.W. and Earles, J.L. (2001) Less really is more for adults learning a miniature artificial language. *J. Mem. Lang.* 44, 250–273
38 Rose, S.A. *et al*. (2001) Visual short-term memory in the first year of life: Capacity and recency effects. *Dev. Psychol.* 37, 539–549
39 Ross-Sheedy, S. *et al*. (2003) The development of visual short-term memory capacity in infants. *Child Dev.* 74, 1807–1822
40 Káldy, Z. and Leslie, A.M. A memory span of one? Object identification in 6.5-month-old infants. *Cognition* (in press)
41 Maurer, D. *et al*. (2002) The many faces of configural processing. *Trends Cogn. Sci.* 6, 255–260
42 DeBoer, T. *et al.* Event-related potentials in developmental populations. In *Event-Related Potentials: A Methods Handbook* (Handy, T., ed.), MIT Press (in press)
43 Reynolds, G.D. and Richards, J.E. Familiarization, attention, and recognition memory in infancy: An ERP and cortical source localization study. *Dev. Psychol.* (in press)
44 Csibra, G. *et al*. (2000) Gamma oscillations and object processing in the infant brain. *Science* 290, 1582–1585
45 Kaufman, J. *et al*. (2003) Representing occluded objects in the human infant brain. *Proc. R. Soc. London B Biol. Sci.* 270(Suppl. 2), S140–S143
46 Deheane-Lambertz, G. *et al*. (2002) Functional neuroimaging of speech perception in infants. *Science* 298, 2013–2015
47 Taga, G. *et al*. (2003) Brain imaging in awake infants by near-infrared optical topography. *Proc. Natl. Acad. Sci. U. S. A.* 100, 10722–10727
48 Baird, A.A. *et al*. (2002) Frontal lobe activation during object permanence: Data from near-infrared spectroscopy. *Neuroimage* 16, 1120–1126
49 Wilcox, T. *et al.* Using near-infrared spectroscopy to assess neural activation during object processing in infants. *J. Biomed. Opt.* (in press)
50 Csibra, G. *et al*. (2004) Near infrared spectroscopy reveals neural activation during face perception in infants and adults. *J. Pediatr. Neurol.* 2, 85–89
51 Peña, M. *et al*. (2003) Sounds and silence: An optical topography study of language recognition at birth. *Proc. Natl. Acad. Sci. U. S. A.* 100, 11702–11705
52 Uttal, W.R. (2003) *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*, MIT Press
53 Culham, J.C. *et al*. (2001) Attention response functions: Characterizing brain areas using fMRI activation during parametric variations in attentional load. *Neuron* 32, 737–745
54 Grill-Spector, H. *et al*. (2004) The fusiform face area subserves face perception, not generic within-category discrimination. *Nat. Neurosci.* 7, 555–562
55 Pascalis, O. *et al*. (2002) Is face processing species-specific during the first year of life? *Science* 296, 1321–1323