Review

Detecting impossible changes in infancy: a three-system account

Su-hua Wang[1] and Renée Baillargeon[2]

[1] Department of Psychology, University of California, Santa Cruz, CA 95064, USA
[2] Department of Psychology, University of Illinois, Champaign, IL 61820, USA

Can infants detect that an object has magically disappeared, broken apart or changed color while briefly hidden? Recent research suggests that infants detect some but not other 'impossible' changes; and that various contextual manipulations can induce infants to detect changes they would not otherwise detect. We present an account that includes three systems: a physical-reasoning, an object-tracking, and an object-representation system. What impossible changes infants detect depends on what object information is included in the physical-reasoning system; this information becomes subject to a principle of persistence, which states that objects can undergo no spontaneous or uncaused change. What contextual manipulations induce infants to detect impossible changes depends on complex interplays between the physical-reasoning system and the object-tracking and object-representation systems.

Change violations and cognitive development

The study of change blindness – the inability to notice changes when they occur during a brief visual disruption – has a long history in the field of infant cognition, and encompasses two broad lines of research. In one, infants are presented with a series of static arrays, and investigators examine whether infants detect salient changes between arrays [1–3]. In the other, infants are presented with physical events in which objects become briefly hidden, and researchers ask whether infants notice 'impossible' changes, or change violations, that occur while the objects are out of sight [4–6]. In this review, we focus on this second line of research.

Investigations of infants' ability to detect change violations have addressed several related questions. First, what change violations can infants detect when shown a single event involving few objects? For example, can infants detect that an object has magically disappeared, broken apart or changed color [7,8]? Second, can infants who fail to detect a change violation be induced to do so through contextual manipulations [9,10]? Third, how does infants' performance deteriorate under more challenging conditions: for example, when the event unfolds rapidly, when the number of objects in the event is increased or when multiple events are shown simultaneously [11,12]? The answers to these questions have proven remarkably intricate, and are helping to shed light on the various cognitive systems that underlie infants' responses of to events.

In this review, we focus on the first two questions (see Box 1 for an overview of the third). Research on the first question – what change violations do infants spontaneously detect? – has helped us formulate an account of the physical-reasoning system of infants [13–15]. Research on the second question – can infants be induced through contextual manipulations to detect change violations they do not spontaneously detect? – is leading us to explore possible links between the physical-reasoning system and two other systems suggested by findings in the adult and infant visual cognition literature: the object-tracking system [16–18], and another system we term the object-representation system [19–22].

Together, these research efforts suggest that at least three different systems – the object-tracking, object-representation and physical-reasoning systems – are needed to fully explain infants' responses to change violations. In what follows, we briefly outline this three-system account, and then review some of the evidence that supports it.

A three-system account

Consider the following situation: infants see a rigid cover standing next to a toy on an apparatus floor; an experimenter lifts the cover, lowers it over the toy and then lifts it again. Do infants detect that the toy has disappeared, or changed size, shape, pattern or color? To answer these questions, we must first consider how infants would represent the event. Below is a description of this process (see Figure 1).

When infants notice the cover and toy, their object-tracking system assigns an index to each object, based on the available spatiotemporal information [16] (because the cover and toy occupy different locations in space, they are readily perceived as separate objects). Each index functions as an index finger or pointer that 'sticks' to its object and tracks it as it moves. In our simple event, the object-tracking system has no difficulty establishing a continuous trace for each object: it knows at all times where each object is located, although the toy is hidden part of the time [18].

As soon as indexes are assigned to the cover and toy, infants' object-representation system begins to build a detailed representation of each object, listing various features. Each object representation is bound to its index [16] so that infants can keep track of the properties of each object. A variety of segregation, recognition and categorization processes can operate on the representations, to include additional information or to highlight particular

Corresponding author: Wang, S.-h. (suhua@ucsc.edu).

1364-6613/$ – see front matter doi:10.1016/j.tics.2007.10.012 Available online 19 December2007

Box 1. Detecting change violations under more challenging conditions

The present review focuses on infants' ability to detect a change violation in a single event involving few objects. As might be expected, the infants' performance often deteriorates under more challenging conditions. Below are two examples.

Increasing the number of objects in the event

In a recent experiment [11], 10.5-month-olds saw two dolls being placed on an apparatus floor. Next, a screen was raised to hide the dolls, and a distractor glided across the apparatus, passing briefly in front of the screen. After the distractor exited the apparatus, the screen was lowered to reveal either one or two dolls. Infants looked reliably longer at the one-doll outcome when the distractor consisted of a single object (four connected beads) but not four separate objects. These results suggest that infants could not form a coherent physical representation of the event when the total number of objects (i.e. two dolls, one screen and four separate beads) exceeded the number of indexes available [16,18,46,47].

Increasing the number of events

In a series of experiments, infants were shown two occlusion events side by side: one toy disappeared behind a first screen and a different toy then disappeared behind a second screen. Results indicated that, at 4 months, infants detected a violation if both screens were lifted simultaneously to reveal the two toys behind the same screen but not in the reverse locations [48]. At 6.5 months, infants detected a violation if the two screens were lifted to reveal the two toys in the reverse locations; however, this success seemed to depend on their keeping track of the toy hidden behind the second screen only [49]. Finally, at 9 months, infants kept track of the toy hidden behind each screen and thus detected a violation if the first screen was lifted to reveal the wrong toy [12]. These results suggest that location binding poses special difficulties for young infants. In time, perhaps with neurological maturation [12,48,49], infants begin to overcome these difficulties, and succeed in binding the information about each event to its proper location.

information [23–27]. For example, infants might use information from previous encounters with the cover (or similar covers) to determine that it is a cover rather than a block.

As the experimenter lifts the cover, infants' physical-reasoning system also becomes involved. The physical-reasoning system is a causal-reasoning system that monitors physical events as they unfold [14,15]. As the event progresses, the physical-reasoning system builds a specialized physical representation of the event. All of the information included in this representation becomes subject to infants' core knowledge [28,29]. This knowledge includes a principle of persistence, which states that an object can undergo no spontaneous or uncaused change in the course of an event: objects are expected to persist, with all of their physical properties, in time and space [14,30].

When building their physical representation, infants first represent the basic information about the event, which comprises both spatiotemporal and identity information. The spatiotemporal information specifies how many objects are involved in the event, and how their arrangement changes as the event unfolds. The identity information provides categorical or ontological information about each object, such as whether it is self- or non-self-propelled [31,32] and whether it is closed or open (e.g. is the object open at the top to form a container, open at the bottom to form a cover, or open at both ends to form a tube? [15,33]).

After infants have represented the basic information about the event, they use this information to categorize the event [15,34–37] (e.g. if the cover is lowered over the toy, infants categorize the event as a covering event; if the cover is lowered in front of the toy, they categorize it as an occlusion event). Infants then tap their knowledge of the event category selected, which lists the variables that have been identified as relevant for predicting outcomes in the category [38–42] (for a description of how variables are identified, see Refs [43,44]). A variable both calls infants' attention to a certain type of information in an event, and provides a causal rule for interpreting this information. For example, in covering events the variable height calls infants' attention to the relative heights of the cover and object, and specifies that the object can become fully hidden under the cover if it is shorter but not taller than the cover.

After infants have determined which variable information they need to include in their physical representation of an event, they gather this information, either directly from the event if the objects are still visible or by querying the object-representation system if not. In either case, variable information is then included in the physical representation of the event and interpreted in accord with the variable rule and core principles. Events that do not unfold as expected are flagged as violations.

As in the object-representation system, the basic and variable information about each object in the physical-reasoning system is bound to its index, so that infants can keep track of the properties of each object. A crucial point here is that the information about an object in the physical-reasoning system might constitute only a small subset of the information about the same object in the object-representation system.

What change violations can infants detect?

Basic and variable change violations

Let us now return to our example: a cover is lowered over a toy and then lifted. What change violations should infants be able to detect (see Box 2)?

Our account predicts that change violations that involve only basic information (or basic change violations) should be detected early, because even young infants would include this basic information in their physical representation of the event. Thus, young infants should detect a violation when the cover is lifted to reveal no toy, and results confirm this prediction [15] (Figure 2).

By contrast, change violations that involve variable information (or variable change violations) should be detected only if infants have included information about the variables in their physical representation of the event. Thus, infants who have identified height, width, shape, pattern and color as covering variables should include information about these variables in their physical representation of the event, and hence should detect changes involving these variables. Conversely, infants who have not yet identified these variables should be blind to all such changes: when the cover is lifted to reveal a toy that differs in height, width, shape, pattern or color from the original toy, infants should give no evidence that they detect the change.

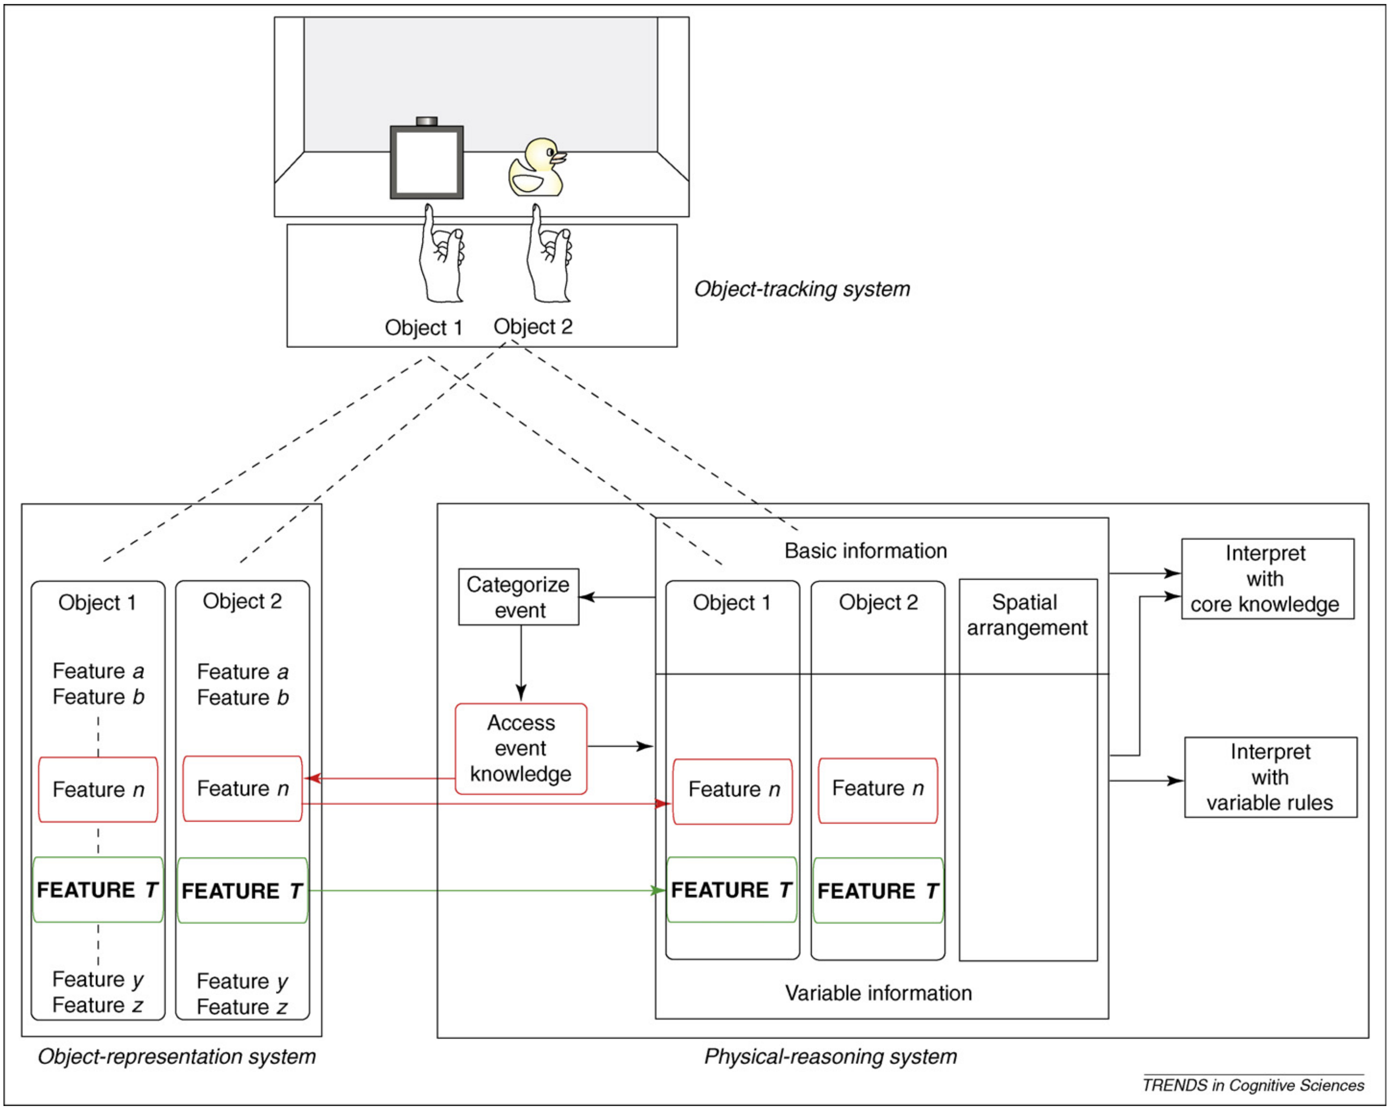

Figure 1. Schematic diagram of three systems involved in infants' representation of a simple physical event. In the example used here, a rigid cover first stands next to a toy. According to the account described in the text, the object-tracking system first assigns an index to each object, and uses these indexes to track the cover and toy as they move. The object-representation system then builds a detailed representation of each object, listing various features (e.g. features *a* to *z* stand for height, width, shape, pattern and so on). As the cover is lifted and lowered over the toy, the physical-reasoning system begins to build a physical representation of the event. First to be represented is the basic information about the event, which includes spatiotemporal information about the number and arrangement of the objects, and identity information about each object (e.g. is the object self-propelled or not, and closed or not?). The basic information is then used to categorize the event. Next, infants access their knowledge of the category selected, which specifies the variables that have been identified as relevant for predicting outcomes in that event category; information about each variable is then included in the physical representation. If the physical-reasoning system cannot obtain information about a variable directly from the event, it can query the object-representation system for it; for example, here information about feature *n* is retrieved from the object-representation system and included in the physical-reasoning system. If a feature has been primed and rendered salient in the object-representation system, information about this feature can be passed on to the physical-reasoning system even if not requested; for example, here information about highlighted FEATURE *T* is passed on to the physical-reasoning system. Any variable information included in the physical-reasoning system is interpreted in accord with the variable rules (if available) and with infants' core knowledge, which includes a principle of persistence.

Evidence for this analysis comes from experiments focusing on the variable height in covering events, which is identified at ~12 months [15,44]. As predicted, 12-month-olds look reliably longer when shown a change violation in which a short object becomes much taller when briefly hidden under a tall cover, whereas 11-month-olds are blind to this change [5] (Figure 3). This blindness does not stem from a fundamental inability to encode, store or maintain information about the relative heights of the cover and object – this information is listed in the object-representation system – but rather from a failure to include the information in the physical-reasoning system.

Event category effects with variable change violations

Because infants learn separately about each event category, and because sometimes months separate the identification of the same variable in different event categories (primarily due to lack of exposure), infants who detect a variable change violation in one event category can fail to detect a similar violation in another category. For example, at 4.5 to 5 months of age, infants detect a shape change if a box changes into a ball when passing behind a narrow screen that is only slightly larger than the box [7]. However, they do not detect a shape change if a clear box containing beads is buried in one location in a sandbox and a mesh ball containing a bell is

Box 2. Superficial changes or change violations?

Most of the research reviewed here uses the violation-of-expectation method, which relies on infants' tendency to look longer at events that violate, as opposed to confirm, their expectations [50]. With change violations, the following interpretational question arises: when infants look reliably longer at a change than at a no-change event, do they do so because they notice a perceptual change or because they implicitly recognize that the change is inconsistent with their physical knowledge? Converging evidence from other conditions and tasks can help address this issue. Below are two examples from experiments with covering events.

At 2.5 to 3 months, infants look reliably longer (i) when a cover is lowered over a toy duck, slid to the side, and then lifted to reveal no duck as opposed to the duck; and (ii) when a cover is lowered over a toy duck, slid behind the left half of a screen slightly taller than the duck, lifted above the screen, moved to the right, lowered behind the right half of the screen, slid past the screen, and finally lifted to reveal the duck as opposed to no duck [15] (see Figure 2). Together, these results make clear that infants are not simply noticing a superficial change from a duck to no duck; rather, they are reasoning about the displacement of the duck and detecting its magical disappearance or appearance.

At 12 months, infants not only look reliably longer when a short object becomes much taller after being briefly hidden under a tall cover [5] (Figure 3), but also search for a tall object under a tall as opposed to a short cover [44], and they look reliably longer when a short cover is lowered over a tall object until it becomes fully hidden [15]. Younger, 11-month-old infants fail all of these tasks [5,15,44]. Together, these results suggest that 12-month-olds expect an object to retain its height when under a cover, and respond accordingly in both violation-of-expectation and action tasks.

then retrieved from the same location [4]. Presumably, infants identify the variable shape in occlusion events before they do so in 'burial' events because they are infrequently exposed to the latter events.

Other experiments examined infants' ability to detect the same change to the same object in different event categories. As predicted, infants detected the change only when they had identified the variable as relevant for the event category (Box 3). Some of these experiments ([5]; Figure 3) involved the variable height, which is identified at ~3.5 months in occlusion events [45], 7.5 months in containment events [36,37], 12 months in covering events [15,44] and 14 months in tube events [15]. Results showed that 11-month-olds detected a change to the height of an object that was briefly hidden behind but not under a cover, and 12.5-month-olds detected a change to the height of an object that was briefly hidden under a cover but not inside a tube (for reviews of additional experiments showing event category effects with variable violations see Refs [14,30]).

Can infants be induced to detect change violations?

Priming effects

According to our three-system account, the physical-reasoning system can query the object-representation system for variable information about the objects in an event. Under usual circumstances, the object-representation system provides only information about queried variables; because variables not yet identified are unlikely to be queried, information about these variables is typically not passed on to the physical-reasoning system. However, what if contextual manipulations were used to render information about an unidentified variable highly salient in the object-representation system? This information might then be passed on to the physical-reasoning system. Once included in the physical representation of the event, the information would become subject to the persistence

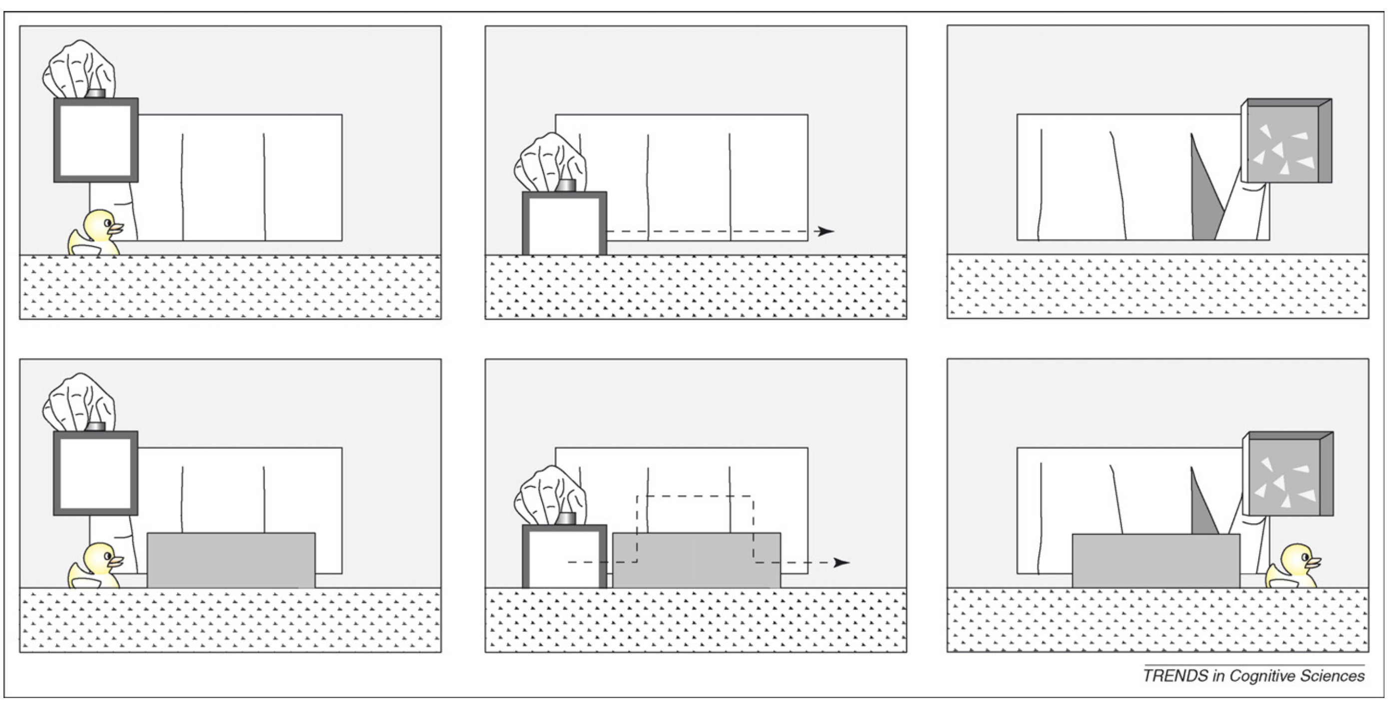

Figure 2. Examples of basic change violations [15]. Top row: A gloved hand holds a cover over a toy duck (the cover is rotated in previous familiarization trials to enable infants to inspect its hollow interior). The hand lowers the cover over the duck, slides the cover to the right, and finally lifts and rotates it. Infants aged 2.5 to 3 months look reliably longer when the duck is absent than present at the end of the event, suggesting that they detect the magical disappearance of the duck. Bottom row: The gloved hand lowers the cover over the duck, and slides the cover behind the left half of a screen slightly taller than the duck. Next, the hand lifts the cover above the screen, moves it to the right, lowers it behind the right half of the screen, slides it past the screen, and finally lifts and rotates it. Infants aged 2.5 to 3 months look reliably longer when the duck is present than absent at the end of the event, suggesting that they detect the magical appearance of the duck (they realize it should have been left behind the left half of the screen).

Box 3. Sparser object representations in the physical-reasoning system

According to the indexing account offered by Leslie, Scholl and their colleagues [16,18], although indexes are assigned based on spatio-temporal information, featural information can then be bound to the indexes, enabling infants to detect violations involving this information. This account has much in common with our own: in particular, both assume that (i) infants' representation of an object in an event can fail to include information about one or more features of the object; and (ii) this featural information can nevertheless be available in a different system (a 'feature map' in the indexing account; the object-representation system in our account).

Nevertheless, there exists at least one crucial difference between the two accounts: whereas the indexing account assumes that all featural information bound to an index is used when reasoning about an event, our account does not. We distinguish between the rich object representations of the object-representation system, and the potentially much sparser object representations of the physical-reasoning system. This distinction is essential to make sense of event category effects: if infants use whatever information is bound to an index when reasoning about an event, it becomes difficult to explain why they can detect a surreptitious change to an object in one event, but fail to detect the same change to the same object in another event.

Why is the physical-reasoning system designed this way? Why does it not simply use the rich representations formed by the object-representation system? One possible answer has to do with learning [14,43,44]. When faced with unexplained contrastive outcomes in an event category (e.g. an object under a cover is sometimes fully and sometimes only partly hidden), the physical-reasoning system searches for and eventually identifies the variable that helps predict these outcomes (e.g. the object becomes fully hidden when it is shorter but not taller than the cover). From that point on, information about the variable (e.g. height) is regularly included when representing events from the category, as demonstrated by the successful performance of infants in violation-of-expectation and action tasks involving the variable [43,44]. If information about all variables were included from the start for all events, infants would have an impossible task of sorting through all of this information. Instead, in each event category, variables are added one by one, as their predictive power becomes clear.

principle of the infants, enabling them immediately to detect change violations involving the information.

Evidence for these speculations comes from seminal experiments by Wilcox and Chapa [10], who primed 7.5-month-olds to detect a color change in an occlusion event (color is not identified as an occlusion variable until ~11.5 months [7]). After receiving three pairs of priming trials in which green cups were used to pound pegs and red cups were used to pour salt, infants detected a change violation when shown a test event in which a green ball became red when passing behind a narrow screen. Our interpretation of these results is that the priming trials rendered the colors green and red highly salient in the object-representation system; as a result, information about the color of the balls was more likely to be passed on to the physical-reasoning system. Once included, the color information became subject to the persistence principle, and the event was flagged as a violation: given the relative sizes of the screen and ball, only one ball could be hidden behind the screen at one time; if the ball was green when it went behind the screen, it could not be red when it came out. Control results supported this interpretation: when the screen was wide enough to hide two balls, infants no longer viewed the event as a change violation (for reviews of additional priming experiments, see Refs [14,30]).

Carry-over effects

According to our three-system account, when watching an event, infants assign indexes to the objects in the event, and use these indexes to keep track of the objects as they move. Recent research has revealed an important function for this tracking process [9]: when infants see a sequence of

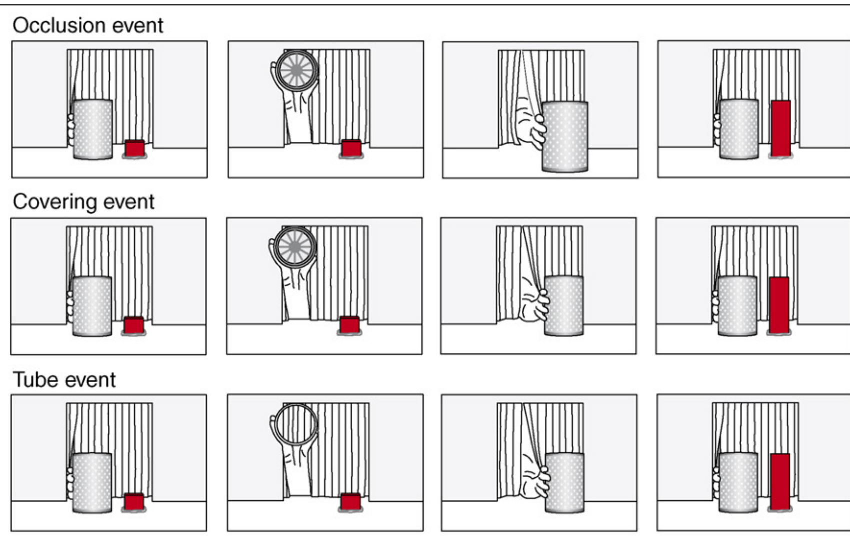

Figure 3. Examples of event category effects [5]. Top row: A gloved hand rotates a cover standing next to a short toy, to enable infants to inspect the hollow interior of the cover. Next, the hand lowers the cover in front of the toy, thus producing an occlusion event. Finally, the hand lifts the cover to reveal a much taller toy. At 11 months, infants detect the surreptitious change to the height of the toy, because they already have identified height as an occlusion variable [35,36,44] and thus include height information in their physical representation of the event. Middle row: In this event, the cover is lowered over the toy, to produce a covering event. Because height is not identified as a covering variable until ~12 months [15,43], 11-month-olds now fail to detect the change to the height of the toy, but 12.5-month-olds succeed in doing so. Bottom row: In this event, the cover is replaced with a tube (identical to the cover with its top removed), which is lowered over the toy. Because height is not identified as a tube variable until ~14 months [15], 12.5-month-olds now fail to detect the change to the height of the toy.

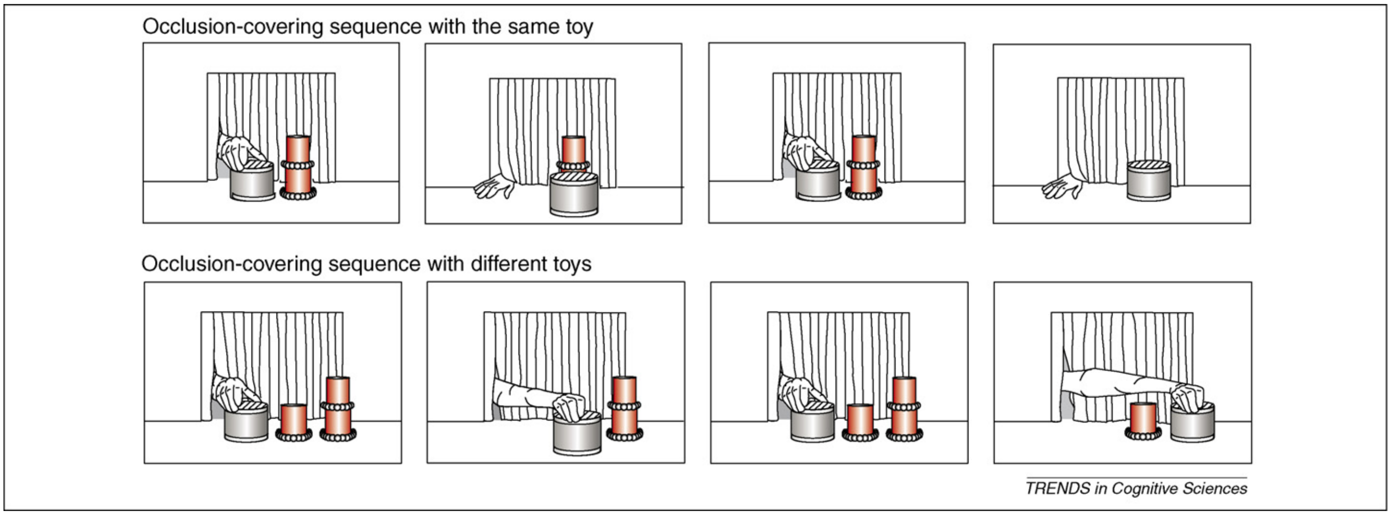

Figure 4. Examples of carry-over effects [9]. Top row: After seeing a short cover being placed in front of a tall toy, 8-month-olds detect a violation if the cover is next lowered over the toy until it becomes fully hidden. If infants see the covering event alone, or if they see the short cover being placed next to, instead of in front of, the toy, they do not detect this violation until ~12 months [9,15]. Bottom row: The same carry-over effect is observed if the short cover is placed in front of a short toy and then lowered over the tall toy, as long as the short and tall toys are greatly similar. If the short and tall toys differ in color, infants no longer detect the violation in the covering event [9].

two distinct events involving the same objects, and the tracking system can unambiguously track the objects from the first to the second event, the variable information included in the physical representation of the first event is carried over *en bloc* to that of the second event. Such a strategy is extremely efficient: why go through the effortful process of gathering variable information about the objects in the second event, when they are clearly the same as in the first event?

The carry-over of variable information is particularly advantageous for infants in situations where they first see an event in which a variable has been identified, followed by an event in which the variable has not yet been identified. Information about the variable is then carried over, fortuitously, from the physical representation of the first event to that of the second event. Once included, the variable information becomes subject to the persistence principle, enabling infants to detect violations they would not otherwise have been able to detect.

The initial evidence for carry-over effects came from experiments [9] (Figure 4) in which we attempted to induce infants to detect an interaction violation involving height information in a covering event (in interaction violations, the respective properties of the objects involved make their interaction impossible). These experiments built on previous findings that the variable height is identified at ~3.5 months in occlusion events, but only at ~12 months in covering events. After watching a short cover being placed in front of a tall toy, 8-month-olds detected a violation when the cover was lowered over the toy until it became fully hidden. According to our account, height information was included in the physical representation of the first, occlusion event; because the object-tracking system could establish continuous traces for the cover and toy from the first to the second event, this height information was then carried over to the physical representation of the second, covering event, enabling infants to detect the violation in the event four months before they would normally have done so. A control experiment supported this interpretation: infants no longer detected the violation in the covering event if the short cover was first placed next to (rather than in front of) the tall toy [9].

In a final experiment [9], the short cover was placed in front of a short toy (fully hiding it) before being lowered over the tall toy. This short toy was either perceptually similar to the tall toy, or different in color. Infants detected the violation in the covering event in the first case but not the latter, pointing to a possible subtle interplay between the object-tracking, object-representation and physical-reasoning systems. (For preliminary experiments showing carry-over effects with change violations, see Ref. [14].)

Concluding remarks

The research reviewed here suggests that infants do not need to learn that objects cannot spontaneously appear or disappear, break apart or coalesce, or change size, shape, pattern or color: from an early age, infants interpret events in accord with a principle of persistence, which deems all such changes impossible. However, whether infants actually succeed in detecting an impossible change in an event depends on whether they have included the necessary information in their physical representation of the event

Box 4. Questions for future research

- What basic identity information (beyond self- or non-self-propelled, and open or closed) do infants obligatorily encode about objects? Do infants show evidence of encoding this information in simple action tasks and in tasks that require detecting change or interaction violations, as suggested by our account?
- We have suggested that the object-representation system can include information about an object that is not included in the physical-reasoning system. Is the reverse ever true? Can the physical-reasoning system include variable information that is not included in the object-representation system, creating a double dissociation?
- What is the time course of priming and of carry-over effects? Furthermore, do these effects have long-term consequences, such as facilitating the identification of variables by infants? If infants are repeatedly induced to include height information in covering events, does it make it easier for them to identify this variable when finally exposed to appropriate observations to do so?

– either spontaneously, or as a result of priming or carry-over manipulations. To make sense of these findings, we have outlined a new account that posits links between three separate systems: an object-tracking, an object-representation and a physical-reasoning system. (For further research questions derived from this account, see Box 4.)

Acknowledgements
The redaction of this review was supported by a Junior Scholar Grant from the Chiang Ching-kuo Foundation and a Special Research Grant from the University of California, Santa Cruz, to S.W., and by a grant from NICHD (HD-21104) to R.B. We thank Jie Li, Amy Needham, Dan Simons and four anonymous reviewers for helpful comments.

References
1 Feigenson, L. *et al.* (2002) Infants' discrimination of number vs. continuous extent. *Cognit. Psychol.* 44, 33–66
2 Huttenlocher, J. *et al.* (2002) Infants and toddlers discriminate amount: are they measuring? *Psychol. Sci.* 13, 244–249
3 Oakes, L.M. *et al.* (2006) Rapid development of feature binding in visual short-term memory. *Psychol. Sci.* 17, 781–787
4 Newcombe, N. *et al.* (1999) Infants' coding of location in continuous space. *Infant Behav. Dev.* 22, 483–510
5 Wang, S. and Baillargeon, R. (2006) Infants' physical knowledge affects their change detection. *Dev. Sci.* 9, 173–181
6 Wilcox, T. and Baillargeon, R. (1998) Object individuation in infancy: the use of featural information in reasoning about occlusion events. *Cognit. Psychol* 37, 97–155
7 Wilcox, T. (1999) Object individuation: infants' use of shape, size, pattern, and color. *Cognition* 72, 125–166
8 Wynn, K. (1992) Addition and subtraction by human infants. *Nature* 358, 749–750
9 Wang, S. and Baillargeon, R. (2005) Inducing infants to detect a physical violation in a single trial. *Psychol. Sci.* 16, 542–549
10 Wilcox, T. and Chapa, C. (2004) Priming infants to attend to color and pattern information in an individuation task. *Cognition* 90, 265–302
11 Cheries, E.W. *et al.* (2006) Interrupting infants' persisting object representations: an object-based limit? *Dev. Sci.* 9, F50–F58
12 Káldy, Z. and Leslie, A.M. (2003) Identification of objects in 9-month-old infants: integrating 'what' and 'where' information. *Dev. Sci.* 6, 360–373
13 Baillargeon, R. *et al.* (2006) Under what conditions do infants detect continuity violations? In *Processes of Change in Brain and Cognitive Development: Attention and Performance XXI* (Munakata, Y. and Johnson, M.H., eds), pp. 163–188, Oxford University Press
14 Baillargeon, R. *et al.* An account of infants' physical reasoning. In *Learning and the Infant Mind* (Woodward, A. and Needham, A., eds), Oxford University Press (in press)
15 Wang, S. *et al.* (2005) Detecting continuity violations in infancy: A new account and new evidence from covering and tube events. *Cognition* 95, 129–173
16 Leslie, A.M. *et al.* (1998) Indexing and the object concept: developing 'what' and 'where' system. *Trends Cogn. Sci.* 2, 10–18
17 Pylyshyn, Z.W. (2001) Visual indexes, preconceptual objects, and situated vision. *Cognition* 80, 127–158
18 Scholl, B.J. and Leslie, A.M. (1999) Explaining the infants' object concept: Beyond the perception/cognition dichotomy. In *What is Cognitive Science?* (Lepore, E. and Pylyshyn, Z., eds), pp. 26–73, Blackwell
19 Angelone, B.L. *et al.* (2003) The roles of representation and comparison failures in change blindness. *Perception* 32, 947–962
20 Kahneman, D. *et al.* (1992) The reviewing of object files: object-specific integration of information. *Cognit. Psychol.* 24, 175–219
21 Hollingworth, A. (2005) The relationship between online visual representation of a scene and long-term scene memory. *J. Exp. Psychol. Learn. Mem. Cogn.* 31, 396–411
22 Mitroff, S.R. *et al.* (2004) Nothing compares 2 views: change blindness results from failures to compare retained information. *Percept. Psychophys.* 66, 1268–1281
23 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
24 Needham, A. (1999) The role of shape in 4-month-old infants' segregation of adjacent objects. *Infant Behav. Dev.* 22, 161–178
25 Needham, A. *et al.* (2005) Infants' formation and use of categories to segregate objects. *Cognition* 94, 215–240
26 Needham, A. *et al.* (2006) Infants' use of category knowledge and object attributes when segregating objects at 8.5 months of age. *Cognit. Psychol.* 53, 345–360
27 Quinn, P.C. (2005) Young infants' categorization of humans versus nonhuman animals: roles for knowledge access and perceptual process. In *Building Object Categories in Developmental Time: 32nd Carnegie Mellon Symposium on Cognition* (Gershkoff-Stowe, L. and Rakison, D., eds), pp. 107–130, Erlbaum
28 Leslie, A.M. (1994) ToMM, ToBY, and agency: Core architecture and domain specificity. In *Mapping the Mind: Domain Specificity in Cognition and Culture* (Hirschfeld, L.A. and Gelman, S.A., eds), pp. 119–148, Cambridge University Press
29 Spelke, E. (1994) Initial knowledge: six suggestions. *Cognition* 50, 431–445
30 Baillargeon, R. Innate ideas revisited: For a principle of persistence in infants' physical reasoning. *Perspect. Psychol. Sci*. (in press)
31 Luo, Y. *et al.* Young infants' reasoning about events involving self- and nonself-propelled objects. *Cognit. Psychol*. (in press)
32 Saxe, R. *et al.* (2005) Secret agents: 10- and 12-month-old infants' inferences about hidden causes. *Psychol. Sci.* 16, 995–1001
33 Hespos, S.J. and Baillargeon, R. (2001) Knowledge about containment events in very young infants. *Cognition* 78, 207–245
34 Aguiar, A. and Baillargeon, R. (2003) Perseverative responding in a violation-of-expectation task in 6.5-month-old infants. *Cognition* 88, 277–316
35 Baillargeon, R. and Wang, S. (2002) Event categorization in infancy. *Trends Cogn. Sci.* 6, 85–93
36 Hespos, S.J. and Baillargeon, R. (2001) Infants' knowledge about occlusion and containment events: a surprising discrepancy. *Psychol. Sci.* 12, 141–147
37 Hespos, S.J. and Baillargeon, R. (2006) Décalage in infants' knowledge about occlusion and containment events: Converging evidence from action tasks. *Cognition* 99, B31–B41
38 Aguiar, A. and Baillargeon, R. (2002) Developments in young infants' reasoning about occluded objects. *Cognit. Psychol.* 45, 267–336
39 Huettel, S.A. and Needham, A. (2000) Effects of balance relations between objects on infants' object segregation. *Dev. Sci.* 3, 415–427
40 Kotovsky, L. and Baillargeon, R. (1998) The development of calibration-based reasoning about collision events in young infants. *Cognition* 67, 311–351
41 Luo, Y. and Baillargeon, R. (2005) When the ordinary seems unexpected: evidence for incremental physical knowledge in young infants. *Cognition* 95, 297–328
42 Wang, S. *et al.* (2003) Should all stationary objects move when hit? Developments in infants' causal and statistical expectations about collision events. *Infant Behav. Dev.* 26, 529–567
43 Wang, S. and Baillargeon, R. Can infants be "taught" to attend to a new physical variable in an event category? The case of height in covering events. *Cognit. Psychol*. (in press)
44 Wang, S. and Kohne, L. Visual experience enhances infants' use of task-relevant information in an action task. *Dev. Psychol*. (In press)
45 Baillargeon, R. and DeVos, J. (1991) Object permanence in young infants: further evidence. *Child Dev.* 62, 1227–1246
46 Feigenson, L. *et al.* (2004) Core systems of number. *Trends Cogn. Sci.* 8, 307–314
47 Feigenson, L. and Carey, S. (2005) On the limits of infants' quantification of small object arrays. *Cognition* 97, 295–313
48 Mareschal, D. and Johnson, M.H. (2003) The "what" and "where" of object representations in infancy. *Cognition* 88, 259–276
49 Káldy, Z. and Leslie, A.M. (2005) A memory span of one? Object identification in 6.5-month-old infants. *Cognition* 97, 153–177
50 Wang, S. *et al.* (2004) Young infants' reasoning about hidden objects: evidence from violation-of-expectation tasks with test trials only. *Cognition* 93, 167–198