Journal of Experimental Psychology: General
2015, Vol. 144, No. 3, 570–580

0096-3445/15/$12.00 http://dx.doi.org/10.1037/a0037750

Visual Perception Involves Event-Type Representations: The Case of Containment Versus Occlusion

Brent Strickland
Yale University and Institut Jean Nicod

Brian J. Scholl
Yale University

Recent infant cognition research suggests that core knowledge involves event-type representations: During perception, the mind automatically categorizes physical events into broad types (e.g., *occlusion* and *containment*), which then guide attention to different properties (e.g., with width processed at a younger age than height in containment events but not occlusion events). We tested whether this aspect of infant cognition also structures adults' visual processing. In 6 experiments, adults had to detect occasional changes in ongoing dynamic displays that depicted repeating occlusion or containment events. Mirroring the developmental progression, change detection was better for width versus height changes in containment events, but no such difference was found for otherwise equivalent occlusion events, even though most observers were not even aware of the subtle occlusion–containment difference. These results suggest for the first time that event-type representations exist and operate automatically and unconsciously as part of the underlying currency of adult visual cognition.

Keywords: event perception, change detection, core knowledge, naive physics

While the light that enters our eyes is continuous, our visual experience of the world is often discrete. We see the world as populated by discrete objects, for example—both object *types* (e.g., seeing some pattern of light as an animal or a vehicle) and object *tokens* (e.g., seeing a particular animal during online perception)—and these discrete object representations serve as the underlying currency of other processes such as memory (e.g., Luck & Vogel, 1997) and attention (e.g., Scholl, 2001). Research also suggests that some particular object types—for example, people and animals—are prioritized in various ways in visual processing (e.g., New, Cosmides, & Tooby, 2007; Thornton, Pinto, & Shiffrar, 1998).

However, our experience of the world is also intrinsically dynamic: We perceive the world not only in terms of objects but also dynamic events. Accordingly, a great deal of recent work has explored the ways in which the mind segments continuous ongoing experience into discrete event tokens (for a summary, see Shipley & Zacks, 2008) and how those token representations then influence other processes such as memory (e.g., Radvansky, Tamplin, & Krawietz, 2010; Swallow, Zacks, & Abrams, 2009) and attention (e.g., Levin & Varakin, 2004; Newtson & Engquist, 1976). Some have also suggested that stereotypical motions of objects (e.g., how a child's swing moves) may be stored as such (e.g., Cavanagh, Labianca, & Thornton, 2001) and that certain specialized types of motion patterns trigger particular percepts (e.g., of biological motion or animacy; e.g., Blake & Shiffrar, 2007; Gao, Newman, & Scholl, 2009). To our knowledge, however, previous work has not explored the idea that the visual system categorizes motion into a privileged set of general event types. Here, inspired by work in infant cognition, we explore the possibility that perception automatically categorizes dynamic visual information into one of a small number of core event types (e.g., occlusion or containment) and that this categorization in turn automatically and unconsciously controls the particular visual features to which we attend.

Event Types in Infant Cognition

For infants, perhaps even more than for adults, there is far too much detail in visual input to fully process, so visual experience must be selective, often via the operation of attention. To which features of an event will an infant's attention be directed? According to a proposal by Renee Baillargeon and her colleagues, infants automatically categorize dynamic visual input into physical event types (e.g., occlusion and containment), which then serve to guide attention to different properties (for reviews, see Baillargeon, 2008; Baillargeon & Wang, 2002).[1] For example, containment might orient infants toward width information, because this variable will be especially relevant to whether an object can be inserted into a potential (vertically oriented) container.

This article was published Online First April 13, 2015.

Brent Strickland, Department of Psychology, Yale University, and Institut Jean Nicod; Brian J. Scholl, Department of Psychology, Yale University.

For helpful conversation and comments on drafts of this article, we thank Renee Baillargeon, Steve Franconeri, Frank Keil, Josh Knobe, Gordan Logan, David Pietraszewski, Laurie Santos, and Phil Wolff.

Correspondence concerning this article should be addressed to Brent J. Strickland, Institut Jean Nicod, UMR 8129, Pavillon Jardin, Ecole Normale Supérieure, 29 rue d'Ulm, F-75005 Paris, France, or to Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT, 06520-8205. E-mail: stricklandbrent@gmail.com or brian.scholl@yale.edu

[1] We use the term *event types* here as a way to contrast types versus tokens in event perception. Baillargeon and colleagues have not used this terminology, referring instead to *event categories*.

Physical event-type representations concern interactions between at least two distinct physical objects. Empirical studies of this possibility have focused on several different potential event types (including covering, collision, and tube events; see Baillargeon, 2008), but the power of this proposal is especially well illustrated in a series of studies that explored infants' attention to the height and width of objects in occlusion versus containment events. In a study of attention to height information (Hespos & Baillargeon, 2001), infants were shown an object that was lowered either behind an occluder or into a container. In some cases, the object was too tall to fully fit behind the occluder or into the container, and looking-time methods were used to assess whether infants noticed this discrepancy. They did so at 4.5 months for occlusion events but not until 7.5 months for containment events, even though these two event types were visually very similar (see also Hespos & Baillargeon, 2006, for converging evidence using a search task). In a related study of attention to width information (Wang, Baillargeon, & Brueckner, 2004), infants were first shown an occluder or a container that was subsequently covered by a screen. An object was then lowered behind the screen, which was then dropped to reveal only the occluder or the container. In some cases, the object was too wide to fully fit behind the occluder or be inserted into the container, and looking-time methods were used to assess whether infants noticed this discrepancy. They did so at 4 months for both occlusion and containment events. These results can be collectively summarized by noting that infants prioritize width over height in containment events but that they treat width and height identically in occlusion events–where the prioritization in such studies is revealed via the ages at which they succeed at taking the relevant variable into account.

Core Knowledge

The type of binding between event types and particular visual properties described above may constitute a form of core knowledge (Spelke, 2000; Spelke & Kinzler, 2007). Abstractly, core knowledge refers to relatively encapsulated and domain-specific mechanisms for representing and processing information about specific domains that play an especially foundational role in people's mental lives, serving as types of primitive representations and underlying many forms of higher level representation and learning (and attending). Purported examples of core knowledge include representations of objects (for a review, see Hood & Santos, 2009), numbers (for a review, see Feigenson, Dehaene, & Spelke, 2004), and even social ingroup–outgroup status (for a review, see Spelke, Bernier, & Skerry, 2013), among others. In each case, the provocative core knowledge claim is that these types of processing are a deep part of human nature rather than being (say) cultural constructs.

More concretely, core knowledge abilities seem to be early emerging (often in infancy), universal (being shared even across otherwise widely diverging cultures), and heavily conserved by natural selection (and so also appearing in nonhuman primates). These features seem to capture the examples mentioned above, for instance (as discussed in the reviews cited in the previous paragraph), but they do not apply to most other examples (contrast representations of political affiliation, derivatives, or astronomy).

Key to the notion of core knowledge is the idea that the core representations serve to guide further learning about the world, and one way they may do this is by guiding attention to specific types of distinctions (e.g., about how objects move, about how the cardinality of different groups compare, or even about how the accents of different speakers differ). In our view, the results of Baillargeon and colleagues reviewed in the previous section may fit comfortably within this framework (though we take care to note that while Baillargeon herself adopts a type of core-knowledge framework for physical reasoning abilities, she thinks that event-type representations are acquired by this system through a process of explanation-based learning; see Baillargeon et al., 2012, for a review). Event types—notably including occlusion and containment—may be a deep part of human nature and may guide attention to distinctions that are of special utility and importance in our experience of the world.

The Current Project: Event Types in Adults' Visual Perception?

Here we ask whether event-type representations may exist and guide attention in adults' visual processing. We ask here about the study of automatic visual processing by way of contrast to higher level cognition. Although it has been notoriously difficult to offer a precise characterization of the differences between perception and cognition, those differences surely exist. In particular, there are powerful generalizations that seem to hold for "perceptual" processes but not "cognitive" ones. Visual perception refers in this context to a family of processes that are relatively automatic and irresistible and that operate without the ability to consciously introspect their nature–while being strongly and directly controlled (unlike other automatic processes of social cognition, say) by specific and subtle features of the visual input itself (for an elaborated characterization, see Scholl & Gao, 2013). Such processes include everything from edge detection and stereopsis to the perception of causality and animacy. In contrast, there are many forms of higher level judgment that are eminently resistible and introspectible and that are not controlled by specific subtle visual details. Here we allude to adults' "automatic" and "perceptual" processing in describing our results for three primary reasons along these lines. To foreshadow, (a) adults' dynamic change detection is driven by differences in especially subtle visual details, (b) this is true even when the observers are completely unaware of those relevant visual details (or even that a difference exists at all), and (c) observers seem ("automatically") constrained by such factors and cannot simply decide what features to readily detect. These features are diagnostic of adults' visual processing in this context (though of course there is no direct way to index the degree to which the parallel results in infants are driven by perception vs. cognition).

On the surface, visual processing in adults may seem relatively unrelated to the study of core knowledge in infant cognition. However, recent work has suggested that these two seemingly different fields may in fact be studying the same underlying representations and constraints. This perspective proposes that core principles and processes in infant cognition are not abandoned after getting development off the ground but rather continue to influence processing in a core fashion throughout life. The core influence that we are interested in here is not the fact that these knowledge structures guide relatively high-level processes such as verbal and/or explicit categorization (e.g., Hespos & Piccin, 2009; Hespos & Spelke, 2004) but instead the intriguing possibility that such representations may guide the operation of relatively automatic and unconscious processes of perception and attention (e.g., Cheries, Mitroff, Wynn, & Scholl, 2008; Mitroff, Scholl, & Wynn, 2004; for a review, see Cheries, Mitroff, Wynn, & Scholl, 2009).

Are physical event types similarly a part of core perceptual processing that operates throughout the lifespan? In the present studies, we explore this question, with two theoretical aims. First, we seek to find out whether such representations function in adult perception at all, since to our knowledge no previous studies of event perception in adults have raised this possibility. Second, by exploring event types in the context of adults' perception, we seek to determine the underlying nature of such representations, as studied in both infants and adults. In particular, the present studies allow us to ask about the degree to which such processing is automatic and occurs intrinsically as a part of visual perception (rather than, e.g., being only a higher level cognitive strategy).

Thus we hypothesize that event types are at their root (a) a particular scheme of categorization for dynamic visual input—of carving up the input such that certain patterns are treated as similar and others as distinct, and (b) a mechanism for automatically prioritizing visual attention to different properties (e.g., the width of an object) as a result of categorization as a particular type (e.g., containment). In what follows, we sometimes speak of *event types* as referring simply to the different possible core categories themselves (e.g., containment, occlusion) and sometimes to the resulting guidance of attention (such that a constitutive feature of the event-type representations is that they do guide attention in this way).

To explore whether event-type representations are part of the currency of adult visual cognition, we sought to ask questions of adults that were analogous to the infant studies of Baillargeon and colleagues. Of course, it would not make sense to use the identical displays and methods, for several reasons. For example, whereas looking times are a reliable indicator of infants' attention, they are much less reliable for adults, whose looking patterns are frequently controlled in a more top-down fashion. Second, adults may often fail to share infants' surprise when viewing seemingly impossible events, either because they immediately grasp the underlying mechanism or because they understand that anything is possible in computer-rendered displays (such as those used in the present study). In any case, our goal is not to replicate the superficial methods of the infant studies but rather to explore the same underlying questions, using the best methods available.

Thus, in order to explore event types in adult perception, we created a novel task in which adults had to detect occasional changes in ongoing dynamic displays that depicted many repeating occlusion or containment events (see Figure 1). In containment events, a rectangular object continuously moved from the edge of the display into a depicted container near the center of the display, and back (see Figure 2A). Occlusion events were identical, except the object repeatedly appeared to move behind the container instead of into it (see Figure 2B). Several of these events were simultaneously present in each 45- to 60-s display, with their motions desynchronized. (These displays are intrinsically dynamic and are difficult to depict in static figures, but animations can be viewed online at http://www.yale.edu/perception/event-types/) Occasionally, a moving object changed either its height or width slightly while it was invisible during the depicted occlusion or containment, and observers had to press a key whenever they detected such a change. Inspired by the theory of event types from infant cognition, we predicted that width

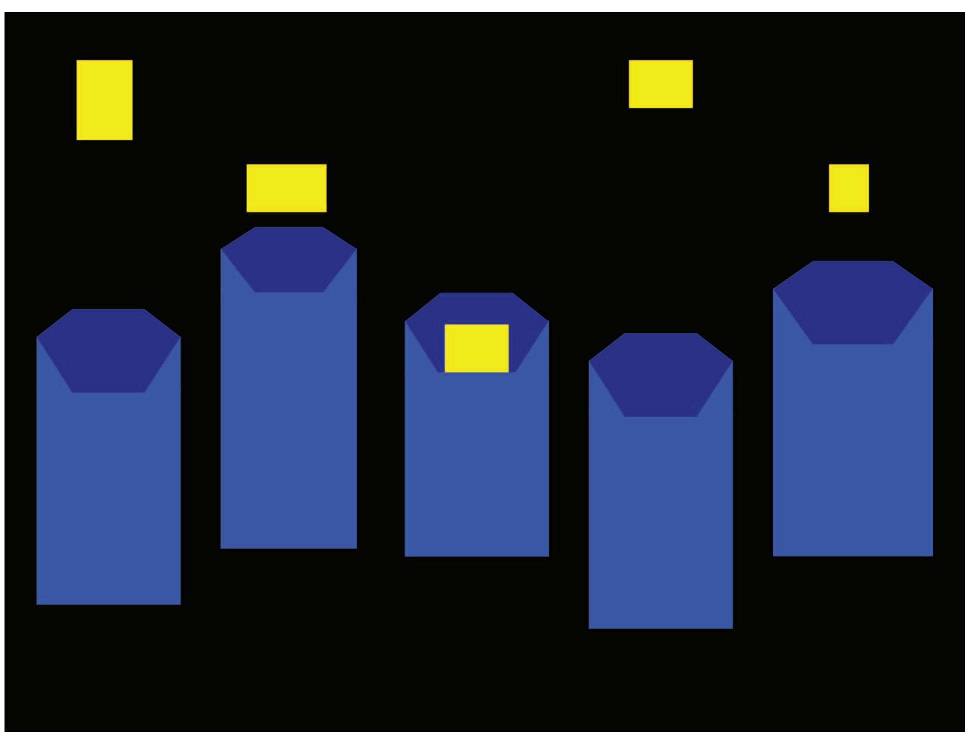

Figure 1. Sample screenshot from a trial in Experiment 1. Each object moved continuously along a vertical path between the top of the screen and the middle of the static container, disappearing either into or behind the depicted container as described in the text. Occasionally, one of the moving objects changed its width or height while invisible, and observers' task was to detect such changes. See the online article for the color version of this figure.

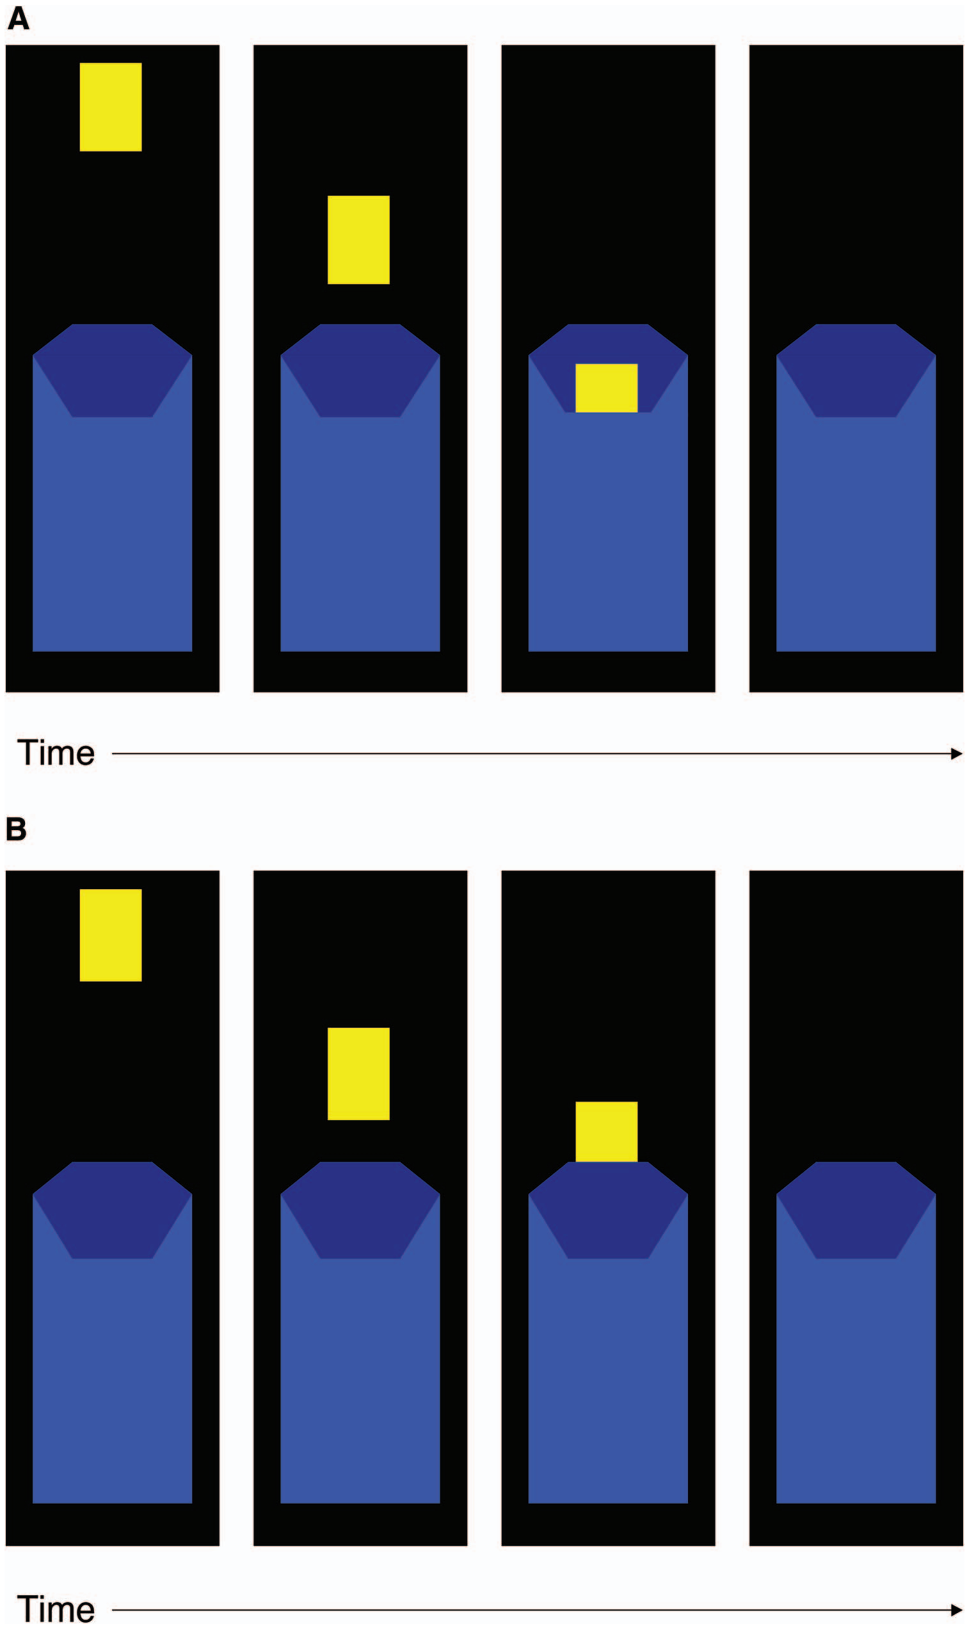

Figure 2. A: Depiction of a containment event from Experiment 1 in which the moving object disappears inside the static container. B: Depiction of an occlusion event from Experiment 1 in which the moving object disappears behind the static container. See the online article for the color version of this figure.

changes would be detected more readily than height changes in containment events, but that no such difference should obtain for highly similar occlusion events.[2]

Experiment 1: Vertical Events

We typically think of containers as vertically oriented, with openings at their tops, and so we first contrasted containment versus occlusion with vertically oriented occluders (as in Figures 1 and 2). Observers were asked to detect changes in shape, with no mention of the different event types.

Method

Participants. Fifteen adults participated in individual 20-min sessions in exchange for a small monetary payment. This sample size was chosen to be in line with many past change detection studies, from our lab and from many others. (This was a mostly haphazard choice, because the dynamic change detection paradigm used here was developed specifically for this study. In particular, given the lack of any previous studies with this exact method, there was no basis for power analyses that would have generated a more precise estimate of the required number of observers.) The data from one observer were not included in the analyses because of a false alarm rate more than 2 standard deviations above the group mean. All observers were naive to the purposes of the experiment.

Apparatus. The experiment was run on a MacBook laptop with an LCD screen using custom software written in Matlab with the Psychophysics Toolbox libraries (Brainard, 1997; Pelli, 1997). Observers sat approximately 51 cm from the screen without head restraint so that the entire display subtended 31.62° × 19.76° of visual angle.

Stimuli. All stimuli were presented within a central black 19.22° square region of the display surrounded by a gray background, shifted up 0.54° from the bottom of the screen. Each display contained five repeating events, each depicting a static container/occluder (henceforth just *container*) near the center of the display, and a rectangle that moved vertically between the top of the display region and the center of the container. All containers in a trial were drawn in either blue or green (randomly chosen), with rectangles correspondingly drawn in either yellow or red. Containers were each drawn in two shades of their color, with a darker polygon near their tops depicting the container opening (as depicted in Figures 1 and 2). Each container had a width chosen randomly between 3.08° and 3.46° and a height chosen randomly between 3.69° and 5.54°. The containers' horizontal positions were evenly spaced throughout the display, and each had its center randomly placed vertically between 8.65° and 14.42° from the top of the display region (see Figure 1).

Each rectangle was horizontally aligned with center of its container and was initially placed out of sight, with its lowest edge one pixel above the top edge of the display region. So that motions were asynchronous, each rectangle began moving after a randomly chosen delay between 0 and 3 s from the beginning of the trial. Each rectangle proceeded to move down to the center of the container (where it was invisible, as described below) and then back up until it again went offscreen at the top of the display region. This cycle repeated for the duration of each trial, randomly chosen between 45 and 60 s. Each time the rectangle disappeared (at either the top or the bottom part of its cycling trajectory)—and at the beginning of the trial—its speed was randomly set between the minimally and maximally extreme values of 6.92°/s and 19.23°/s. Each rectangle always disappeared near the center of the display by either occlusion (with its leading edge disappearing as it contacted the initial contours of the container, as depicted in Figure 2B) or containment (with its leading edge disappearing only as it contacted the brighter contours depicting the nearer edge of the container's opening, as depicted in Figure 2A). Containment versus occlusion was assigned randomly for each rectangle but remained constant throughout each trial.

After all objects had appeared, a change was triggered after a randomly chosen delay between 3 and 6 s. Once a change was triggered, a random container was selected, and its corresponding rectangle changed its size the next time that it reached the center of the container (which, on average, required an extra delay of 1.08 s). Changes then continued to be triggered in the display every 3 to 6 s (randomly chosen, starting with the previous trigger), with an average interchange latency of 4.45 s. There were an average of 9.8 changes per trial, and the final change had to occur at least 2 s before the end of a trial.

The distributions of height and width changes were always identical, and change magnitudes were always determined irrespective of containment versus occlusion conditions. Each rectangle was initially either wide (with its width randomly between 27% and 32% of the corresponding container's height, and its height randomly between 27% and 32% of the corresponding container's width) or tall (with these values reversed). When a rectangle changed its aspect ratio, either its height or width (randomly chosen) increased by a factor of 1.8 (i.e., 180%) or decreased (randomly chosen) by a factor of 1/1.8 (i.e., 55.6%). Across all changes in a trial, these values varied between the maximally and minimally extreme values of 207% and 37% of the original size of the object at the beginning of the trial. The moving rectangle's width never exceeded the width of its container as it occupied at most 97% of the container's width (as defined by the distance at the widest point of the opening of the container). The moving object's height also never exceeded the height of its container, and it occupied at most 71% of the containers height (as defined by the distance from the bottom of the container to the bottom of its opening). The average magnitude across all changes for both event types and change types was virtually identical at approximately 0.72° in each of the four conditions. Similarly, the minimum and

[2] Note that the displays used in many of the infant studies reviewed above used impossible interactions between physical objects such as solidity violations (to which infants looked longer), but the relevant events involved in our displays were simply the (unlikely, unexplained, but perhaps not impossible) spontaneous changes of the object sizes themselves. In particular, our displays involved no impossible interactions between objects as in many infancy studies (e.g., an object seeming to pass through another object). The questions being asked are nevertheless analogous, however, because the use of impossible physical interactions here (as in most such studies) is not critical to the theoretical questions being asked about event types but is rather a part of the methodology used to collect meaningful data from infant looking time in the first place. In Baillargeon's theory, for example, event types are thought to be active all the time, and not only in the weird or impossible physical interactions that were often used to measure such representations with infants. (Indeed, some analogous results have been observed in infant studies that do not use impossible events, per se; e.g., Wang & Baillargeon, 2006; Wang & Mitroff, 2009)

maximum change magnitudes were equated across conditions at approximately 0.39° and 1.37°, respectively.

Procedure. Observers were instructed to attend broadly to the display and to make a speeded response via a keypress whenever they detected the occurrence of an aspect-ratio change. These instructions were conveyed without mentioning containment or occlusion. Responses made within 1.75 s of an aspect-ratio change (starting at the first moment in which the change could have been detected in principle) were coded as hits, and others were coded as false alarms. After two practice trials (the results of which were not recorded), observers completed 10 experimental trials.

During postexperimental debriefing (in this experiment and also in Experiments 2 and 6), observers were asked whether they had noticed the difference between occlusion and containment events, and the analyses were conducted only over that subset of observers (always more than 60%) who reported not having noticed the difference. We limited the analyses in this way in an attempt to isolate relatively automatic perceptual processing, as opposed to conscious strategies, because conscious strategies could not readily operate over distinctions that observers did not notice in the first place. (Nevertheless, in this and all subsequent experiments, the patterns of results did not change qualitatively when all observers were analyzed: In all cases, the patterns of significance were qualitatively the same, and the significant effects reported below were quantitatively more significant when all observers were included.)

Results

The changes—an average of 97.14 per observer—were analyzed without regard for the trial from which they came. In response to explicit questioning during the debriefing, nine of the 14 observers denied noticing the difference between occlusion and containment during the experiment. False alarm rates were collected for each observer, along with hit rates for each of the four categories (height changes in containment events, width changes in containment events, height changes in occlusion events, width changes in occlusion events). The average false alarm rate was 32.70%, and the average overall hit rate was 49.58%.[3] In planned comparisons, we compared change detection for height versus width in the two event types. (We did not similarly compare occlusion versus containment for the two spatial dimensions, because no such results could speak directly to the questions being asked.) For containment events, width changes were detected marginally more often than height changes (57.31% vs. 45.93%), $t(8) = 2.29$, $p = .052$, $\eta_p^2 = .395$, but no such difference obtained for occlusion events (48.34% vs. 43.40%), $t(8) = 0.86$, $p = .417$, $\eta_p^2 = .084$.

Discussion

The results of this first experiment are consistent with the possibility that adults' visual perception involves event-type representations that guide selective attention: Here, as in the infant studies that directly motivated this work (Hespos & Baillargeon, 2001; Wang et al., 2004), viewing containment events appears to have led observers to selectively attend to width information, whereas viewing occlusion events did not result in any relative prioritization of width or height.

Experiment 2: Horizontal Events

We hypothesized, following the initial infant work (see Section 3 of Wang et al., 2004), that width was prioritized over height during containment events in Experiment 1 because only width was relevant to the determination of whether the object could initially be inserted into the container's opening–and thus whether the containment relation could hold at all. (Height, in contrast, was of only secondary relevance, determining not whether containment could hold at all but just whether the entire object could fit within the container.) In contrast, both variables were equally relevant to occlusion. If this is the right sort of explanation, then it makes the strong concrete prediction (and one that has not been tested with infants, to our knowledge) that the relative prioritization of width and height should flip for containment events that involve horizontal containers (where *width* and *height* are defined here in an absolute sense—i.e., relative to the overall display—rather than being defined relative to the container). We tested this possibility in the current experiment, and the relevant conditions can be depicted simply by turning Figures 1 and 2 by 90°.

Method

This experiment was identical to Experiment 1 except as noted here. Fourteen new observers participated. Given that Experiment 1 produced a robust result with 14 initially tested observers (after one outlier was removed), this same sample size was used in this experiment (and in Experiments 3–5), in part to enable cross-experiment comparisons. Observers viewed displays that were identical to those in Experiment 1 but with the contents of the central display region rotated by 90°. (This change was made in the program's underlying code but is equivalent to simply rotating the monitor itself by 90°.)

Results and Discussion

The changes—an average of 99.75 per observer—were analyzed without regard for the trial from which they came, and via the same planned comparisons as in Experiment 1. In response to explicit questioning during the debriefing, it was again the case that nine of the 14 new observers denied noticing the difference between occlusion and containment during the experiment. For those participants unaware of the difference between experimental conditions, planned contrasts revealed that height changes in containment events (i.e., height relative to the display) were now detected more often than were width changes in containment events (52.58% vs. 42.50%), $t(8) = 2.64$, $p = .030$, $\eta_p^2 = .466$, but no such difference obtained for occlusion events (40.74% vs. 37.27%), $t(8) = 1.11$, $p = .300$, $\eta_p^2 = .133$.

Superficially, this experiment produced the opposite pattern of results from Experiment 1, with height now prioritized over width during containment events. Functionally, however, this experiment fully replicated Experiment 1, with the prioritized variable during containment again being the one that was relevant to whether the rectangle could be inserted into the opening of the container.

[3] Note that it is not possible to compute signal detection measures using this design, because it is not possible to localize a false alarm to any particular category, as is possible with hits and misses. This issue is explored directly in Experiment 6.

Experiment 3: Event Types Versus Shading Contrast

To depict containers in our simple animations, we necessarily had to depict the inner and outer surfaces of the container differently, and we did so via distinct shading of the same hue, as might be produced by simple lighting differences (as in Figures 1 and 2). This meant, though, that the rectangles were disappearing behind surfaces with different shades in containment versus occlusion. Could those shading differences alone have driven our results? We tested this possibility by directly contrasting change detection with uniform surfaces of each of the two possible shades, as depicted in Figure 3.

Method

This experiment was identical to Experiment 2 (again involving horizontal trajectories) except as noted here. Fourteen new observers (the same sample size as in the previous experiments) participated. Each container in the current experiment was replaced with a single uniform rectangle of the same size (see Figure 3), randomly assigned to either be the darker shade (corresponding to the nearer surface in Experiments 1 and 2) or the lighter shade (corresponding to the more distant surface in Experiments 1 and 2).

Results and Discussion

The changes—an average of 98.89 per observer—were analyzed without regard for the trial from which they came, and via planned comparisons analogous to those in Experiment 2. The average false alarm rate was 38.08%, and the average overall hit rate was 40.62%. Hit rates for height versus width changes did not differ for either light surfaces (39.45% vs. 39.51%), $t(13) = 0.02$, $p = .987$, $\eta_p^2 < .001$, or dark surfaces (42.19% vs. 38.32%), $t(13) = 0.73$, $p = .476$, $\eta_p^2 = .04$. These results rule out the possibility that the containment versus occlusion contrast could be explained in terms of simple shading.

Experiment 4: Event Types Versus Disappearance Duration

The basic physics and optics of occlusion versus containment in Experiments 1 and 2 required another confounded difference: In containment events, the moving rectangle stayed hidden for a slightly shorter amount of time than it did in the occlusion events. Here, we directly tested whether this factor influenced change detection by matching the temporal dynamics of Experiments 1 and 2 but testing only occlusion behind uniform surfaces of various sizes.

Method

This experiment was identical to Experiment 2 except as noted here. Fourteen new observers (the same sample size as in the previous experiments) participated. Each container in the current experiment was replaced with a single uniform surface whose height (relative to the display as a whole) was the same as in Experiment 2 but whose display-relative width was randomly

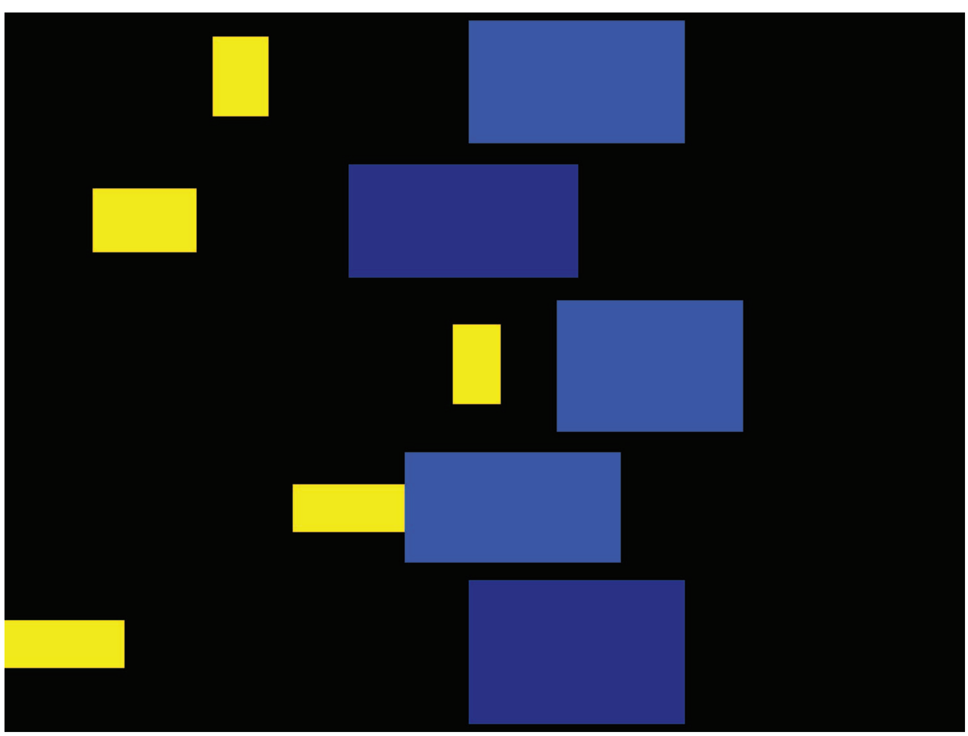

Figure 3. Screenshot of a sample trial from Experiment 3. Objects moved horizontally between the left of the screen and the middle of a static occluder. Each occluder's shade was matched to either the nearer surface or the more distant surface of the containers from Experiments 1 and 2. See the online article for the color version of this figure.

assigned to either be equal to the width of the front surface of the container from Experiment 2 (small) or the width of the of the container as a whole (large).

Results and Discussion

The changes—an average of 99.24 per observer—were analyzed without regard for the trial from which they came, and via planned comparisons analogous to those used in Experiment 2. The average false alarm rate was 38.29%, and the average overall hit rate was 50.34%. Hit rates for height versus width changes did not differ for either small surfaces (51.41% vs. 46.47%), $t(13) = 1.22$, $p = .244$, $\eta_p^2 = .103$, or large surfaces (53.01% vs. 49.50%), $t(13) = 0.86$, $p = .404$, $\eta_p^2 = .054$. These results rule out the possibility that the containment versus occlusion contrast from Experiments 1 and 2 reflects simple timing differences.

Experiment 5: Obtuse Versus Reflex Angles

The basic optics of occlusion versus containment in Experiments 1 and 2 also necessitated a third confounded difference: In containment events, the edge of the surface from behind which the moving rectangle appeared (i.e., the front of the container) was defined by a series of obtuse angles (i.e., greater than 90° but less than 180°). However, the relevant edges for occlusion events (i.e., the back of the container) were defined by a series of reflex angles (i.e., greater than 180° but less than 360°). So, perhaps these different types of notches influenced change detection; in particular, perhaps changes in the container-insertion-relevant dimension were easier to detect in the presence of the obtuse angles (e.g., because the surrounding peaks provide a type of ruler by which the moving rectangles' widths can be compared).[4] In this experiment, we directly tested whether this difference influenced change detection by creating occluding surfaces whose top edge was defined either by a series of obtuse or reflex angles that matched the relevant angles in Experiments 1 and 2 (as depicted in Figure 4). All events were occlusion events.

Method

This experiment was identical to Experiment 2 except as noted here. Fourteen new observers (the same sample size as in the previous experiments) participated. Each containment event was replaced with by an occluding surface whose top edge was defined by obtuse angles. This was achieved simply by removing the back of the container, leaving only its front (such that it now no longer appeared to be a container but instead only an oddly shaped occluder). Each occlusion event was replaced by an occluding surface with only reflex angles on its top edge. This was achieved by changing the color of the back of the container to match its front, thus turning the apparent container into a uniform occluding surface.

Results and Discussion

The changes—an average of 99 per observer—were analyzed without regard for the trial from which they came, and via planned comparisons analogous to those used in Experiment 2. The average false alarm rate was 38.05%, and the average overall hit rate was 45.94%. Hit rates for height versus width changes did not differ for the obtuse-edge events (47.54% vs. 45.22%), $t(13) = 0.79$, $p = .446$, $\eta_p^2 = .045$, but height changes were detected more readily than width changes for the reflex-edge events (49.79% vs. 41.07%), $t(13) = 2.51$, $p = .026$, $\eta_p^2 = .326$.

Critically, these results go in the opposite direction from the effects of the relevant conditions in Experiments 1 and 2. In those experiments, the container-insertion-relevant dimension was detected better (than the container-insertion-irrelevant dimension) in the containment events but not in the occlusion events. Here, in contrast, that same dimension (which was height, given the horizontal orientation of the events) was detected better in the occlusion-matched condition (i.e., the reflex-edge events), but there was no difference in the containment-matched condition (i.e., the obtuse-edge events).

This reversal of height–width prioritization is an unexpected and potentially interesting finding that could receive further empirical attention. One possible explanation is that observers actually categorized the reflex-angle events (unlike the obtuse-angle events) as instances of containment (with containment openings rendered invisible due to the apparent fully frontal orientation of the shapes). Another possibility is this reflects a completely independent effect whereby variability along the axis orthogonal to motion direction is more easily detected (as suggested by Liverence & Scholl, 2011). Regardless of which of these (or any other such) possibilities is true, however, we stress here that these results rule out the possibility that the effects observed in Experiments 1 and 2 could be explained by differences in the types of angles from behind which the moving rectangle reappeared. If anything, they instead suggest that the basic geometry of the containment and occlusion events should have biased the results against our original hypothesis.

Experiment 6: Sensitivity Versus Bias

Our primary goal in these studies is to explore the differing patterns of attentional prioritization that may result from viewing different event types, and such patterns could be realized in principle by enhanced sensitivity or by an increased bias to see such information (even when not present), where both possibilities are consistent with differences in hit rates. Both possibilities would be interesting and would constitute a new kind of attentional prioritization in the study of adult visual cognition. However, an interpretation in terms of increased sensitivity would be most consistent with our motivation in at least two ways. First, enhanced sensitivity would most strongly support the idea that event types are operating relatively automatically in perceptual processing (as opposed to influencing responses, postperceptually). Second, enhanced sensitivity would better match the motivation and interpretation of our methodology in terms of actual change detection, per se.

It was not possible to untangle sensitivity and bias in Experiments 1 and 2, because false alarms could not be isolated to any particular category. For example, when an observer mistakenly pressed a key to indicate a detected change when in fact there was none, both containment and occlusion events would still be actively occurring simultaneously, and so the false alarm could not be assigned to any particular event type. In contrast, the present

[4] We thank an anonymous reviewer for suggesting this possibility.

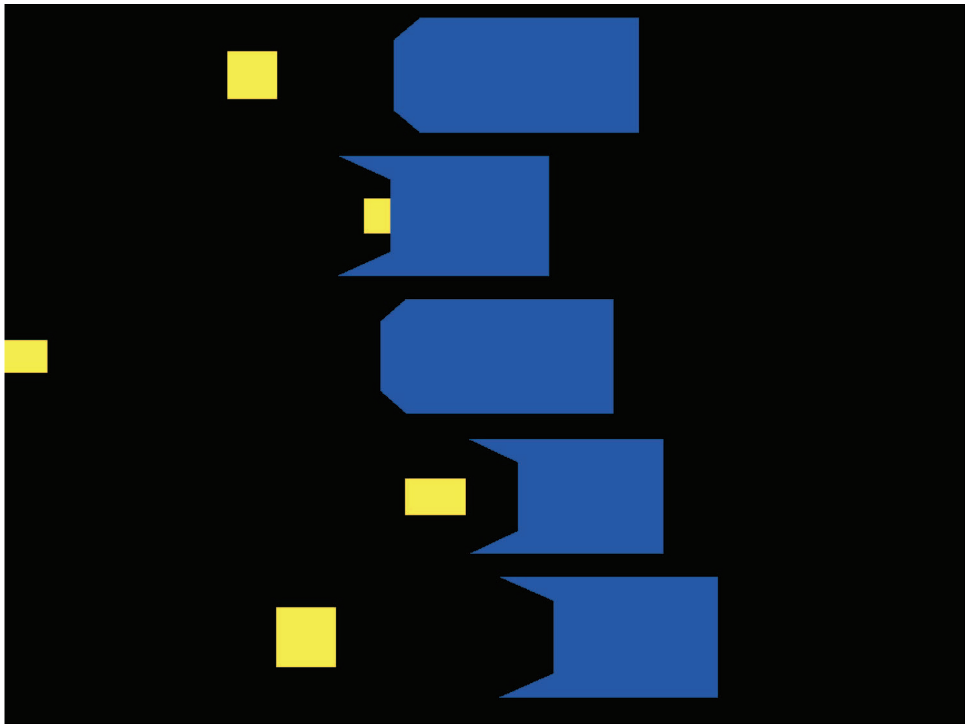

Figure 4. Screenshot of a sample trial from Experiment 5. Objects moved behind occluders, whose ends were composed of a series of either reflex angles or obtuse angles, matched to the structure of either the fronts or the backs of the containers in Experiment 2. See the online article for the color version of this figure.

experiment used the same events as in Experiment 2, but now included a blocked design: In any given trial, all of the depicted events were of one event type (either containment or occlusion), and all of the changes were of one spatial dimension (either height or width).

Method

This experiment was identical to Experiment 2 except as noted here. Twenty-eight new observers participated. This sample size was chosen to precisely double that used in Experiments 1–5, in part due to the concern that the blocked design (described below) could facilitate other types of strategies that would weaken the effect magnitudes (e.g., leading observers to attend in a more focused way, compared with when the nature of the manipulations on a given trial could not be expected, as in the previous studies). This seemed safe, because (unlike the other experiments) there was no possibility of comparing results across experiments, given the substantial shift in methods. The data from one observer were not included in the analyses because of a hit rate more than 2 standard deviations below the group mean. All trials were blocked so that only a single change type and event type occurred on any given trial. Each trial type (height changes with containment, width changes with containment, height changes with occlusion, width changes with occlusion) occurred three times, yielding a total of twelve 45- to 60-s trials.

Results and Discussion

The changes—an average of 117.82 per observer—were analyzed without regard for the individual trial from which they came, and via planned comparisons analogous to those used in Experiment 2. In response to explicit questioning during the debriefing, 20 of the 27 observers denied noticing the difference between occlusion and containment during the experiment. Analyses of the hit rates alone replicated the results of Experiment 2, now in a blocked design. Planned contrasts revealed that height changes in (horizontal) containment events were detected marginally more often than width changes in containment events (49.39% vs. 43.58%), $t(19) = 2.06$, $p = .054$, $\eta_p^2 = .182$, but no such difference obtained for occlusion events (43.87% vs. 43.80%), $t(19) = 0.03$, $p = .979$, $\eta_p^2 = .003$.

To distinguish sensitivity from bias, d' was calculated separately for trials of each of the four categories by subtracting the z score of each participant's false alarm rate from the z score of their hit rate, with the resulting mean values depicted in Figure 5. For containment events, d' was significantly higher for display-relative height changes than display-relative width changes (0.44 vs. 0.07), $t(19) = 2.52$, $p = .021$, $\eta_p^2 = .25$, whereas no such difference obtained for occlusion events (0.10 vs. 0.14), $t(19) = 0.37$, $p = .714$, $\eta_p^2 = .007$—with these two differences themselves yielding a reliable interaction, $F(1, 19) = 7.05$, $p = .016$, $\eta_p^2 = .27$.

These results indicate that the difference in change detection performance between containment and occlusion events derives from true differences in the ability to detect changes (as opposed to biased responding), consistent with the possibility that these differences reflect relatively automatic perceptual processing. Furthermore, isolating sensitivity in this way makes the nature of this effect especially clear: As is starkly illustrated in Figure 5, all of the action in the differences is driven only by increased prioriti-

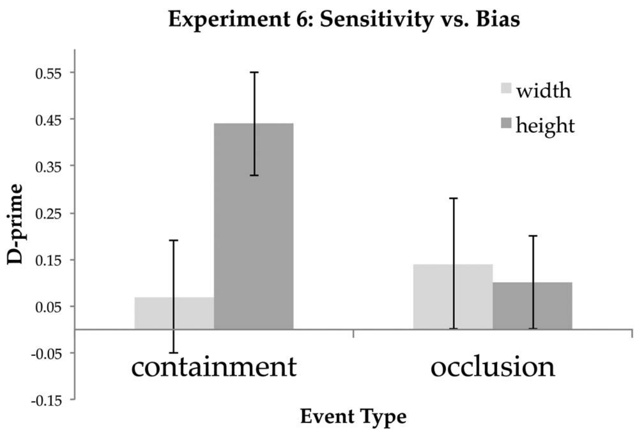

Figure 5. The *d′* values associated with each condition of the blocked design used in Experiment 6. Error bars depict standard errors of the mean.

zation of the dimension that was especially relevant to determining whether the object could be inserted into the opening of the container (i.e., height in this study, because it involved horizontal trajectories).

General Discussion

The experiments presented here demonstrate for the first time that adult visual perception involves event-type representations, wherein certain dynamic events (e.g., containment) automatically prioritize certain questions (e.g., "Could an object even begin to be inserted into the container's opening?"), which in turn prioritize attention to (and memory for) particular properties of objects (e.g., the width of an object that may potentially move into a vertically oriented container). In particular, when contrasting height versus width changes in a dynamic change-detection paradigm, we observed that changes to the container-insertion-relevant dimension were detected more readily for containment events but not for occlusion events. This effect was demonstrated with containers that were oriented both vertically (in Experiment 1) and horizontally (in Experiments 2 and 6) and could not be explained by appeal to subtle differences in contrast (Experiment 3), timing (Experiment 4), or geometry (Experiment 5).

Automaticity

The attentional prioritization that we observed could be because of width (i.e., the container-insertion-relevant dimension) being prioritized globally during containment events or because of it being focused on initially during each event, as that dimension would initially determine whether containment could begin to occur at all (with height becoming relevant only later, to determine whether the object might disappear entirely into the container). In either case, it is especially notable that this prioritization occurred over and over, in rapidly oscillating events. This suggests that the prioritization was triggered by temporally local features of the display in a relatively automatic and data-driven way, because the prioritization occurs even after observers surely know that the same events will be repeating at length.

In addition, and beyond the use of a paradigm adapted from visual cognition research, two other specific features of these results suggest that the attentional prioritization observed here reflects relatively automatic visual processing rather than higher level conceptual strategies or decisions. First, a signal detection analysis (in Experiment 6) directly revealed that the prioritization of the insertion-relevant dimension in containment events reflects enhanced sensitivity to those changes per se. Second, we always limited our analyses in each of the relevant experiments to that majority of observers who had not noticed the containment–occlusion difference (as assessed via careful debriefing). Thus, our results could not reflect an explicit decision to focus more on containment versus occlusion, given that these observers were not aware of this difference in the first place. These results collectively implicate a relatively automatic perceptual phenomenon—perhaps a core part of how vision naturally operates—rather than deliberate strategies.

Core Knowledge in Adulthood

The contrast between occlusion versus containment was initially explored in experiments with infants by Baillargeon and her colleagues (e.g., Hespos & Baillargeon, 2001; Wang et al., 2004) that led to a subsequent theoretical framework based on what we have referred to as *event types*: Infants are seen to categorize event tokens into type representations, which subsequently lead to the prioritization of particular visual features of the token (see Baillargeon, 2008; Baillargeon & Wang, 2002). This mechanism may help to direct infants' attention to the most relevant features of the world (all of the time, and not only when viewing seemingly impossible events) and, in turn, seems to embody a type of core knowledge about how the world is structured.[5]

Whereas core knowledge is typically considered in developmental contexts (e.g., Spelke, 2000; Spelke & Kinzler, 2007) as a tool to help drive progress in infants' reasoning, the current results suggest that event types also continue to serve a core function that continues to shape visual experience in adults. In particular, the fact that such biases persist in adult visual cognition implies that they do not merely get development off the ground but they also help structure adults' mature experience of the world. In this way, the theoretical contribution of the present results is twofold: First, the studies clarify our understanding of the nature of event-type representations in both infants and adults, suggesting in particular that one of their primary functions is to help in the efficient allocation of limited visual processing resources. Second, beyond such developmental links, our results demonstrate the existence of a new type of representation in the adult visual system. In particular, these results suggest that perception automatically categorizes dynamic visual input into one of a small number of core event types (e.g., occlusion or containment) and that this categorization (even when completely unconscious, based on unnoticed cues) in turn controls the particular subtle visual features to which we reflexively attend.

[5] Our interpretation of the results from Experiments 1, 2, and 6 also makes the prediction that infants would notice display-relative height violations earlier in development than display-relative width violations in horizontal containment events, but, to our knowledge, this prediction has yet to be tested.

References

Baillargeon, R. (2008). Innate ideas revisited: For a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science, 3,* 2–13. doi:10.1111/j.1745-6916.2008.00056.x

Baillargeon, R., Stavans, M., Wu, D., Gertner, R., Setoh, P., Kittredge, A. K., & Bernard, A. (2012). Object individuation and physical reasoning in infancy: An integrative account. *Language Learning and Development, 8,* 4–46. doi:10.1080/15475441.2012.630610

Baillargeon, R., & Wang, S. (2002). Event categorization in infancy. *Trends in Cognitive Sciences, 6,* 85–93. doi:10.1016/S1364-6613(00)01836-2

Blake, R., & Shiffrar, M. (2007). Perception of human motion. *Annual Review of Psychology, 58,* 47–73. doi:10.1146/annurev.psych.57.102904.190152

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436. doi:10.1163/156856897X00357

Cavanagh, P., Labianca, A. T., & Thornton, I. M. (2001). Attention-based visual routines: Sprites. *Cognition, 80,* 47–60. doi:10.1016/S0010-0277(00)00153-0

Cheries, E. W., Mitroff, S. R., Wynn, K., & Scholl, B. J. (2008). Cohesion as a principle of object persistence in infancy. *Developmental Science, 11,* 427–432. doi:10.1111/j.1467-7687.2008.00687.x

Cheries, E. W., Mitroff, S. R., Wynn, K., & Scholl, B. J. (2009). Do the same principles constrain persisting object representations in infant cognition and adult perception? The cases of continuity and cohesion. In B. Hood & L. Santos (Eds.), *The origins of object knowledge* (pp. 107–134). Oxford, United Kingdom: Oxford University Press. doi:10.1093/acprof:oso/9780199216895.003.0005

Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences, 8,* 307–314. doi:10.1016/j.tics.2004.05.002

Gao, T., Newman, G. E., & Scholl, B. J. (2009). The psychophysics of chasing: A case study in the perception of animacy. *Cognitive Psychology, 59,* 154–179. doi:10.1016/j.cogpsych.2009.03.001

Hespos, S. J., & Baillargeon, R. (2001). Infants' knowledge about occlusion and containment events: A surprising discrepancy. *Psychological Science, 12,* 141–147. doi:10.1111/1467-9280.00324

Hespos, S. J., & Baillargeon, R. (2006). Décalage in infants' reasoning about occlusion and containment events: Converging evidence from action tasks. *Cognition, 99,* 31–41. doi:10.1016/j.cognition.2005.01.010

Hespos, S. J., & Piccin, T. (2009). To generalize or not to generalize: Spatial categories are influenced by physical attributes and language. *Developmental Science, 12,* 88–95. doi:10.1111/j.1467-7687.2008.00749.x

Hespos, S. J., & Spelke, E. S. (2004, July 22). Conceptual precursors to language. *Nature, 430,* 453–456. doi:10.1038/nature02634

Hood, B., & Santos, L. (Eds.). (2009). *The origins of object knowledge*. Oxford, United Kingdom: Oxford University Press. doi:10.1093/acprof:oso/9780199216895.001.0001

Levin, D. T., & Varakin, D. A. (2004). No pause for a brief disruption: Failures of visual awareness during ongoing events. *Consciousness and Cognition, 13,* 363–372. doi:10.1016/j.concog.2003.12.001

Liverence, B. M., & Scholl, B. J. (2011, November). *Selective inhibition of change detection along the axis of motion: A case study of perception compensating for its own limitations*. Paper presented at the annual Object Perception Attention and Memory meeting, Seattle, WA.

Luck, S. J., & Vogel, E. K. (1997, November 20). The capacity of visual working memory for features and conjunctions. *Nature, 390,* 279–281. doi:10.1038/36846

Mitroff, S. R., Scholl, B. J., & Wynn, K. (2004). Divide and conquer: How object files adapt when an object splits into two. *Psychological Science, 15,* 420–425. doi:10.1111/j.0956-7976.2004.00695.x

New, J., Cosmides, L., & Tooby, J. (2007). Category-specific attention for animals reflects ancestral priorities, not expertise. *Proceedings of the National Academy of Sciences, USA, 104,* 16598–16603. doi:10.1073/pnas.0703913104

Newtson, D., & Engquist, G. (1976). The perceptual organization of ongoing behavior. *Journal of Experimental Social Psychology, 12,* 436–450. doi:10.1016/0022-1031(76)90076-7

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. doi:10.1163/156856897X00366

Radvansky, G. A., Tamplin, A. K., & Krawietz, S. A. (2010). Walking through doorways causes forgetting: Environmental integration. *Psychonomic Bulletin & Review, 17,* 900–904. doi:10.3758/PBR.17.6.900

Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition, 80,* 1–46. doi:10.1016/S0010-0277(00)00152-9

Scholl, B. J., & Gao, T. (2013). Perceiving animacy and intentionality: Visual processing or higher-level judgment? In M. D. Rutherford & V. A. Kuhlmeier (Eds.), *Social perception: Detection and interpretation of animacy, agency, and intention* (pp. 197–230). Cambridge, MA: MIT Press. doi:10.7551/mitpress/9780262019279.003.0009

Shipley, T. F., & Zacks, J. M. (Eds.). (2008). *Understanding events: From perception to action*. New York, NY: Oxford University Press. doi:10.1093/acprof:oso/9780195188370.001.0001

Spelke, E. S. (2000). Core knowledge. *American Psychologist, 55,* 1233–1243. doi:10.1037/0003-066X.55.11.1233

Spelke, E. S., Bernier, E. P., & Skerry, A. E. (2013). Core social cognition. In M. Banaji & S. Gelman (Eds.), *Navigating the social world: What infants, children, and other species can teach us* (pp. 11–16). New York, NY: Oxford University Press. doi:10.1093/acprof:oso/9780199890712.003.0003

Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science, 10,* 89–96. doi:10.1111/j.1467-7687.2007.00569.x

Swallow, K. M., Zacks, J. M., & Abrams, R. A. (2009). Event boundaries in perception affect memory encoding and updating. *Journal of Experimental Psychology: General, 138,* 236–257. doi:10.1037/a0015631

Thornton, I. M., Pinto, J., & Shiffrar, M. (1998). The visual perception of human locomotion. *Cognitive Neuropsychology, 15,* 535–552. doi:10.1080/026432998381014

Wang, S., & Baillargeon, R. (2005). Inducing infants to detect a physical violation in a single trial. *Psychological Science, 16,* 542–549. doi:10.1111/j.0956-7976.2005.01572.x

Wang, S., & Baillargeon, R. (2006). Infants' physical knowledge affects their change detection. *Developmental Science, 9,* 173–181. doi:10.1111/j.1467-7687.2006.00477.x

Wang, S., Baillargeon, R., & Brueckner, L. (2004). Young infants' reasoning about hidden objects: Evidence from violation-of-expectation tasks with test trials only. *Cognition, 93,* 167–198. doi:10.1016/j.cognition.2003.09.012

Wang, S., & Mitroff, S. R. (2009). Preserved visual representations despite change blindness in infants. *Developmental Science, 12,* 681–687. doi:10.1111/j.1467-7687.2008.00800.x

Received June 5, 2012
Revision received July 22, 2014
Accepted July 24, 2014 ■