

Effects of grouping and attention on the perception of causality

HOON CHOI and BRIAN J. SCHOLL
Yale University, New Haven, Connecticut

Beyond perceiving patterns of motion in simple dynamic displays, we can also perceive higher level properties, such as *causality*, as when we see one object collide with another object. Although causality is a seemingly high-level property, its perception—like the perception of faces or speech—often appears to be automatic, irresistible, and driven by highly constrained and stimulus-driven rules. Here, in an exploration of such rules, we demonstrate that perceptual grouping and attention can influence the both perception of causality in ambiguous displays. We first report several types of grouping effects, based on connectedness, proximity, and common motion. We further suggest that such grouping effects are mediated by the allocation of attention, and we directly demonstrate that causal perception can be strengthened or attenuated on the basis of where observers are attending, independent of fixation. Like Michotte, we find that the perception of causality is mediated by strict visual rules. Beyond Michotte, we find that these rules operate not only over discrete objects, but also over perceptual groups, constrained by the allocation of attention.

Long traditions in both philosophy and psychology have concerned themselves with the origin of our notions of cause and effect, precisely because these notions seem to lie at the root of nearly all aspects of reasoning. Hume (1740/1960, 1748/1977) argued that our notions of cause and effect must arise from repeated experiences and noticed correlations, since no mark of causation can be directly perceived from any possible sensory evidence. His hallmark case, discussed at length (especially in the *Enquiry Concerning Human Understanding*), involved the physical causality at work in simple interactions such as one billiard ball colliding with another. He suggested that there is nothing intrinsic in such an event to suggest a causal relation: “Motion in the second Billiard-ball is a quite distinct event from motion in the first; nor is there any thing in the one to suggest the smallest hint of the other” (1748/1977, p. 18).

This associationist view was directly attacked in the early half of the 20th century by the Belgian psychologist Albert Michotte, who in his landmark book *The Perception of Causality* (1946/1963) argued that our minds are, in fact, directly wired to perceive physical causality, even upon first encountering it in billiard-ball-like displays. Michotte constructed displays in which a single disk (A) moves toward a stationary disk (B) until they are adjacent, at which point A stops and B starts moving along the same

path (see Figure 1). We perceive this type of event as a *launch*, or a collision: A smashes into B, causing its motion.¹

Following Michotte, many types of evidence have been marshaled by psychologists in support of the view that the perception of causality is an automatic and possibly even innate process (see Scholl & Tremoulet, 2000, for a recent review). For example, many researchers (including Michotte) have emphasized that whether or not causality is perceived in simple dynamic displays seems to be determined by particular low-level features of those displays, but not by higher level beliefs and intentions (e.g., Michotte, 1946/1963; Scholl & Nakayama, 2002), and other work has shown that causal perception and causal judgment can even pull in opposite directions (Schlottmann & Shanks, 1992). Furthermore, the perception of causality is typically automatic and even irresistible in those displays in which low-level details support it. Other researchers have focused on the universality of causal perception, demonstrating that it occurs across cultures (e.g., Morris & Peng, 1994) and even in very young infants (e.g., Leslie & Keeble, 1987). Thus it seems (not without controversy; see e.g., Schlottmann, 2000; White, 1995) that the *perception of causality* might be aptly named; despite the fact that causation seems like such a higher level cognitive property, there may in fact be perceptual processes that are involved in inferring the causal structure of the world on the basis of spatiotemporal object trajectories, just as there are processes involved in inferring the physical structure of the world based on the basis of more typical visual properties, such as color and shape.

The Principles Underlying the Perception of Causality

Setting aside questions about whether the perception of causality is truly best seen as visual processing, and whether

B.J.S. was supported by NSF Grant BCS-0132444 and by NIMH Grant 1-R03-MH63808-01. For helpful conversation and/or comments on earlier drafts, we thank Marvin Chun, Valerie Kuhlmeier, Alan Leslie, Steve Mitroff, and three anonymous referees. In addition, we thank Vanya Pasheva for assistance with data collection. Correspondence concerning this article should be addressed to B. J. Scholl or H. Choi, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205 (e-mail : brian.scholl@yale.edu, h.choi@yale.edu).




Figure 1. Depiction of a standard causal launching event, wherein one disk is seen to collide with a second disk, *causing* it to move. This event is presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion.

it is innate or modular, many other investigators have focused simply on trying to determine the nature of the *rules* that drive causal perception. This project also began with Michotte (1946/1963), who—in the course of over 100 experiments in the English translation of his landmark book—worked out the details of how perceptual causality is impacted by the absolute and relative speeds of the objects, the distances and directions in which they travel (both before and after “impact”), various types of spatial and temporal gaps, and many other variables. More recent investigations have continued this project, exploring how the perception of causality is affected by different types of motion (e.g., apparent motion; Gordon, Day, & Stecher, 1990), the use of more than two interacting objects (e.g., the *tool effect*; Michotte & Thinès, 1963/1991), other types of spatiotemporal gaps (e.g., Schlottmann & Anderson, 1993), other modalities (e.g., Guski & Troje, 2003), and many other factors (e.g., Boyle, 1960; Gemelli & Cappellini, 1958; Hubbard, Blessum, & Ruppel, 2001; Kruschke & Fragassi, 1996; Natsoulas, 1961; Schlottmann, Allen, Linderoth, & Heskett, 2002; Schlottmann & Shanks, 1992; Weir, 1978; White, in press; White & Milne, 1997, 1999; Yela, 1952).

Causal Capture

All of the studies discussed above explored the impact on perceptual causality of various properties of the objects actually involved in the putatively causal events. In contrast, Scholl and Nakayama (2002; see also Scholl & Nakayama, 2004) demonstrated a contextual effect on the perception of causality, wherein the perceived causal status of a given event (involving objects A and B) was systematically affected by behavior of the other objects (C and D) in a separate event. In these displays, A still initially moved toward a stationary B; eventually, A stopped, and B started moving along the same trajectory, as in a typical launching event. Here, however, A kept moving until A and B were fully overlapped, after which B began its move-

ment (see Figure 2). Percepts from this display were ambiguous: although observers sometimes perceived causal launching (despite the overlap), they often perceived noncausal *passing*, wherein a single moving object was seen to pass right over (or through) a second stationary object that always remained in the center of the screen.

A striking effect emerged in these experiments, however, when a second event—involving two additional objects, C and D—was added to the display (Figure 3). These additional objects behaved as a canonical launch (without overlapping, as in Figure 1). This additional event was added to the ambiguous display described above: C and D moved in a plane below A and B (so that both events were always seen as involving distinct objects), and the two events were spatiotemporally synchronized, so that C and D became adjacent at the same time that A and B fully overlapped (see Figure 3). Despite the addition of this new *context* event involving C and D, observers still reported their percepts of the causal status of only the full-overlap test event, involving A and B. However, with the context event visible, observers were no longer able to perceive A and B as noncausal passing; rather, the unambiguously causal nature of C and D effectively *captured* A and B, so that the test event was also seen as a causal launch. This *causal capture* effect was one of the first demonstrations that the perception of causality is determined not only by the nature of the objects involved in the event to be judged, but also by information from other distinct objects and events in other parts of the display.

Scholl and Nakayama (2002, 2004) explained this effect by appealing to the tendency of the visual system to avoid coincidences when inferring the structure of the world from visual images (e.g., Knill & Richards, 1996; Marr, 1982): The spatiotemporal synchronization of the two events is seen as nonaccidental, and the visual system thus assumes that the ambiguous test event must also be a launching event of the same type as the unambiguous context event.²




Figure 2. Depiction of the full-overlap event. Although this event is sometimes seen as causal launching (despite the overlap), observers often perceive it as *passing*, wherein one moving object is seen to pass right through another stationary object (swapping colors). This event is presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion.




Figure 3. Depiction of the causal capture phenomenon (Scholl & Nakayama, 2002), presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion. Observers judge the causal status of the upper full-overlap event (which, in isolation, is typically seen as noncausal passing). In the presence of the temporally synchronized unambiguous launching context event, the full-overlap event is now also perceived as causal launching.

The Present Experiments: Roles of Grouping and Attention

In presenting the causal capture effect, Scholl and Nakayama (2002) placed great weight on the *causal* nature of the context event: they adverted to perceptual grouping between the two events (in particular, between A and C, and also between B and D), but they suggested that this grouping could not be the full explanation, since no causal capture was observed with various other types of context events, such as a single moving object or a sudden flash. However, it remains possible that other types of context objects and events—even those not perceived as causal launches in isolation—could still impact the perception of the critical ambiguous test event.

Here, we explore this question in general terms: How can various types of perceptual grouping affect the perception of causality? In Experiments 1, 2, and 3, we demonstrate various types of grouping effects on the perception of causality. We began in Experiment 1 (“Causal Dumbbells”) by showing that simple types of connectedness can impact how displays are parsed into individuals and how such parses can correspondingly affect whether ambiguous displays are perceived as causal. In Experiment 2 (“Basic Grouping Effects”), we explored grouping via good continuation, proximity, and common motion, showing that the addition of even a single object to an ambiguous display can cause observers to see either launching or passing, depending on precisely how the additional object is grouped over time with the original objects. Experiment 3 (“Reverse Causal Capture”) further demonstrated the power of perceptual grouping by showing that, in some cases, grouping can even attenuate the perception of causality in an otherwise-*unambiguous* launching event. We interpret such effects in terms of visual attention, and in Experiments 4 and 5, we directly demonstrate that the operation of attention alone can influence the perception of causality.

GENERAL METHOD

Because many of the experiments reported here employed similar methods, we first discuss some of the procedures common to all of them.

Participants

Between 10 and 12 naive Yale students participated in each experiment for course credit or a small monetary payment. All the ob-

servers had normal or corrected-to-normal visual acuity and normal color vision.

Apparatus

All the displays were presented on a Macintosh iMac computer. The observers were positioned approximately 50 cm from the monitor, without head restraint, so that the display subtended approximately $34.2^\circ \times 26^\circ$ of visual angle. The display refreshed at 95 Hz, and motion was always perceptually smooth. All displays were constructed using Macromedia Director animation software and were presented to the observers within an experimental protocol run by MacStim (Darby, 2002). All experiments were conducted in a darkened room.

Stimuli

Each trial involved one or more visual events, each consisting of one or two objects. All objects were small colored disks, each subtending 1.72° , drawn on a black background. The disks were always brightly colored—red, green, or blue. Motion was always in the horizontal plane, since the perception of causality is weaker in other orientations (Michotte, 1946/1963). There were three canonical event types: the ambiguous *full-overlap* event, the unambiguous *launching* event, and an intermediate *partial-overlap* event.

The ambiguous full-overlap event (see Figure 2) proceeded as follows. The initial display consisted of two disks (one red, one green) on the same horizontal plane. The disks were typically vertically centered in the display, though this varied by experiment. One of the disks (A) started out near either the right or the left edge of the display, so that its most extreme edge was 6.79° from the display border. The second disk B was initially drawn near the center of the display. After 333 msec, A began to move at $14.76^\circ/\text{sec}$ toward the center of the display. When the two objects became fully overlapped (with A typically drawn on top of B), A stopped moving, and B instantly started moving at the same speed toward the other edge of the display (stopping when its most extreme edge was 7.31° from the display border). This entire ambiguous test event typically lasted 1.3 sec.

The unambiguous launching event (Figure 1) was identical to the ambiguous full-overlap event, except that Disk A stopped moving, and Disk B started moving, as soon as the two disks became adjacent near the center of the display. The intermediate partial-overlap event (see Figure 8A) was similar, except that Disk A stopped moving, and Disk B started moving, when the two disks were 67% overlapped.

Procedure

On each trial, the observers simply viewed the display and reported, via a keypress, whether they perceived the critical test event as a causal launch (in which one moving object was seen to cause the motion of the other object) or as a noncausal pass (in which one moving object simply traversed the entire length of the display, changing colors, and passing over another stationary object, which also changed colors). Before completing any test trials, the observers first completed a small number of practice trials, the data from

which were not recorded. During these practice trials, the observers were familiarized with the particular conditions employed in that experiment, until they could reliably perceive both launching and passing. The dependent variable in each of the experiments reported here is the percentage of causal launching that was perceived in various conditions—a measure which turns out to be greatly and systematically affected by various stimulus manipulations.

EXPERIMENT 1 Causal Dumbbells

To begin exploring the role of grouping in causal perception, we first manipulated what is perhaps the most basic type of grouping: actual physical connectedness. Prior studies have suggested that connectedness can induce a particularly strong form of perceptual grouping. Palmer and Rock (1994), for example, showed that connectedness could overcome even the powerful effect of proximity. In their demonstration (see Figure 4), observers tended to perceive elements that were connected to each other as being grouped, even when proximity itself supported a conflicting grouping.

In order to explore the effect of grouping by connectedness, we made a new type of stimulus based on the causal capture animation (Scholl & Nakayama, 2002). Recall that the basic causal capture effect (Figure 3) involves the influence of an unambiguous contextual launch event on an ambiguous full-overlap test event. Our baseline *unconnected condition* was constructed by weakening the causal capture display in two ways. First, the test and context events were spatially separated so that the influence of the unambiguous causal context event was attenuated. Second, the test and the context events remained temporally synchronized, but the grouping by common motion was weakened by having the two events proceed in opposite directions (see Figure 5A): for example, the test event might proceed left to right, whereas the context event proceeded right to left. In this baseline stimulus, causal capture does not reliably occur.

In the most important manipulation for this experiment, we constructed a second display—the *dumbbell condition*—which was identical to the unconnected baseline display, except that the two launched disks in each event (called B and D in the initial descriptions above) were always physically connected by a single-pixel line, so that they looked like a dumbbell. As the event proceeded, this line would then rotate, keeping the two disks connected throughout their motion (Figure 5B). The critical question was whether this simple grouping manipulation would cause the ambiguous full-overlap test event to also be perceived as a causal launch, because of its physical connection to the corresponding disk in the unambiguous launch event. Two other conditions were also tested. The *gap condition* was identical to the dumbbell condition, except that the connecting lines never actually contacted the disks—that is, a small gap was always present at each end (Figure 5C). This condition constituted an initial test of whether an effect in the dumbbell condition was driven by con-




Figure 4. A standard demonstration of grouping via connectedness. (A) The disks perceptually group into pairs via proximity. (B) The disks group via connectedness into dumbbells, overriding proximity.

nectedness per se or by a more general form of grouping. Finally, the *removal condition* was identical to the dumbbell condition up until the moment of full overlap in the test event, at which point the connecting line suddenly disappeared for the remainder of the event. This condition allowed us to test the temporal dynamics of grouping by connectedness in this context and to begin to determine precisely *when* grouping plays a role.

Collectively, these four conditions constituted a first test of whether a simple form of perceptual grouping by connectedness could influence the perception of causality.

Method

Twelve observers participated. Each trial involved two visual events: the ambiguous full-overlap test event and the unambiguous launching context event (as described in the General Method section). In both events, the launching disk was bright red, and the launched disk was bright green. The uppermost point of the disks in the test event was 7.40° below the upper display border. The context event was positioned below the test event, so that there was always 5.25° of vertical blank space between the nearest edges of the objects in the different events. The two events moved in opposite directions: In the test event, the launching disk (A) appeared near the left edge of the screen and moved toward the right (eventually overlapping B in the center); meanwhile, in the context event, the launching disk (C) appeared near the right edge of the screen and moved toward the left (eventually reaching D in the center). Despite moving in opposite directions, the two events were always temporally synchronized: A fully overlapped B in the test event at exactly the same moment that C and D in the context event were adjacent.

There were four conditions, as was described above. The unconnected baseline condition worked exactly as just described. The dumbbell condition was identical, except that a 1-pixel green line was always drawn connecting the two green disks in the different events, throughout their motion. The gap condition was identical to the dumbbell condition, except that the connecting line never actually contacted the disks but, rather, was separated from the nearest disk contour by 0.47° . Finally, the removal condition was identical to the dumbbell condition up until the moment of full overlap in the test event, at which point the connecting line suddenly disappeared for the remainder of the event. The observers completed 20 trials in each of the four conditions, for a total of 80 trials; all the trials were presented in a random order, unblocked.

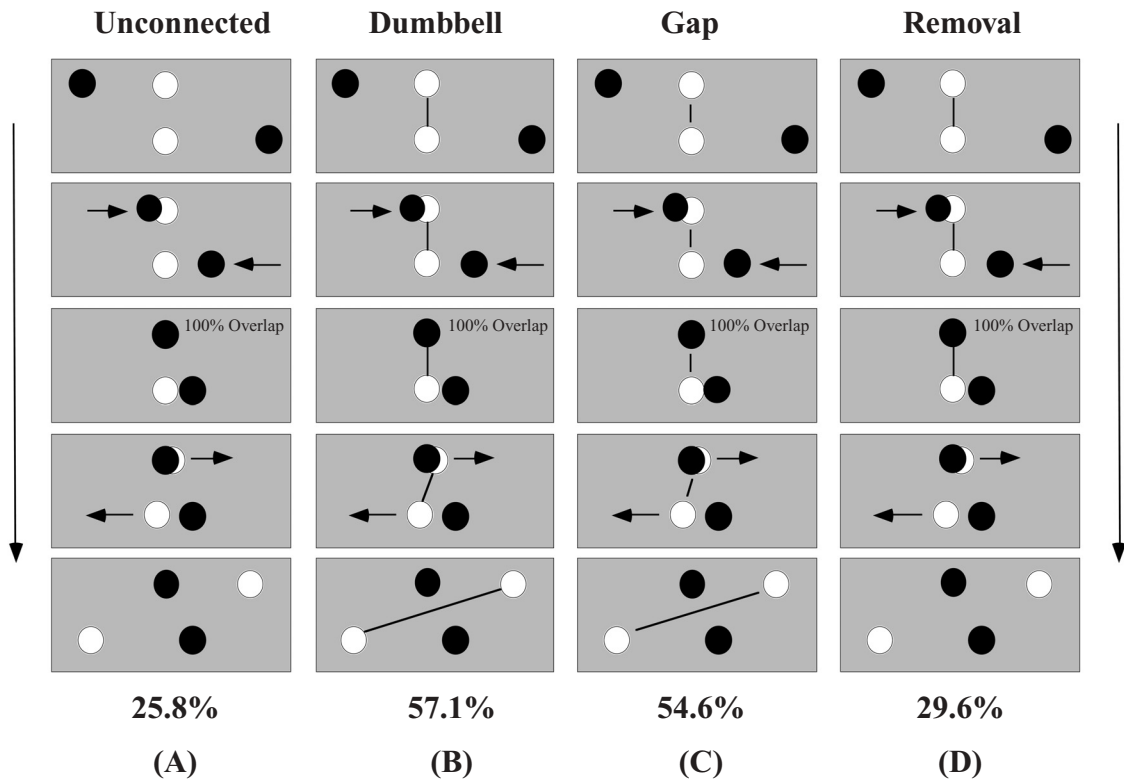


Figure 5. Depictions of the four conditions used in Experiment 1 and the resulting percentages of perceived causal launching. All actual disks were either red or green, as is described in the text. Arrows indicate motion. Each event is presented as a sequence of static frames, with time moving down. The observers always reported their percept of the causal status of the upper event. Animations of these and all other conditions reported in this article can be viewed with a Web browser at <http://www.yale.edu/perception/causal-grouping/>.

Results

The percentage of trials in each condition that were perceived as causal launching are shown at the bottom of each column in Figure 5. A single-factor repeated measures analysis of variance (ANOVA) revealed a significant effect of condition [$F(3,33) = 15.69, p < .001$], and a critical planned comparison revealed that causal launching was perceived more than twice as often in the dumbbell condition (57.1%) as in the unconnected condition [25.8%; $t(11) = 4.34, p = .001$]. Additional planned comparisons revealed no difference in the amount of causal launching percepts between the gap and dumbbell conditions [$t(11) = 1.20, p = .256$] or between the removal and the unconnected conditions [$t(11) = 1.57, p = .145$].

Discussion

These results confirm that physical connectedness between objects in different events can greatly influence the perception of causality. In a weakened causal capture display, the observers tended not to perceive the test event as causal but did perceive causal launching more than twice as often when the unambiguously launched object in the context event was physically connected to the analogous object in the test event. In essence, the added line grouped

the two events, strengthening the causal capture even in this weakened display.

Given that the addition of the line has a pronounced effect on causal perception, we can also ask *when* it has this effect. We might intuitively predict that it is only the presence of the line leading up to the moment of the putative *impact* that drives the effect; after all, the distinction between launching and passing critically concerns this moment in time (the *launch!*), and it is not clear how information *after* this moment could affect causal perception. Nevertheless, no effect of the line was observed in the removal condition, when the line disappeared at the moment of impact. This suggests that the visual system may be inferring the existence of causality not just on a moment-by-moment basis, but by temporally integrating over a larger window of time (cf. Watanabe & Shimojo, 2001). Further experiments that carefully vary precisely when the line is added and removed will be required to determine the nature of this temporal integration.

Beyond concerns of temporal integration, the essential point of this experiment remains the effect of physical connectedness on causal perception. It is tempting to see this as an effect of segmentation: Perhaps the addition of the connecting line caused the visual system to treat the

two launched disks as a single complex object. However, the fact that causality was perceived just as often in the gap condition as in the dumbbell condition suggests that the effect might actually have been driven by a more general form of perceptual grouping. This possibility was explored directly in the following experiment.

EXPERIMENT 2A Basic Grouping Effects

Both the previous experiment and the original causal capture project (Scholl & Nakayama, 2002) explored the rules that mediate the perception of causality by determining which types of context events were able to influence the perception of the ambiguous full-overlap test event. The causal capture phenomenon demonstrates that multiple events can interact in this way, and Experiment 1 suggested that at least one simple form of grouping between the items in two events is sufficient to drive causal capture. But are multiple events necessary to influence the perception of causality in this way? Because Experiment 1 suggested a role for perceptual grouping, we also explored the ways in which *individual* context objects might influence the perception of causality. Unlike previous cases,

the context objects in this experiment do not form unambiguous causal events themselves; there are context *objects*, but there are no interactive context *events*. Nevertheless, grouping between these context elements and the objects in the ambiguous test event can influence the perception of causality in systematic ways.

Our first test of such direct grouping effects employed grouping by good continuation and common motion. We used the same full-overlap test event as that in the previous experiment. In the *baseline condition*, this event appeared in isolation. In the initial display of the other conditions, however, three additional disks were added near the center of the screen, which had the same size and color as Disk B in the test event (i.e., the disk that could be launched when the event was seen as causal). The additional three disks were arranged in a vertical column directly above B, so that they appeared to form a group of four identical disks in a single column (see the top panels of Figure 6).

Three test conditions varied how these objects moved during the motion of the test event. In the *same-motion condition* (Figure 6A), the additional context disks always stayed aligned with B, remaining stationary during A's initial motion and then moving along with B after the overlap in the test event. In the *no-motion condition* (Figure 6B),

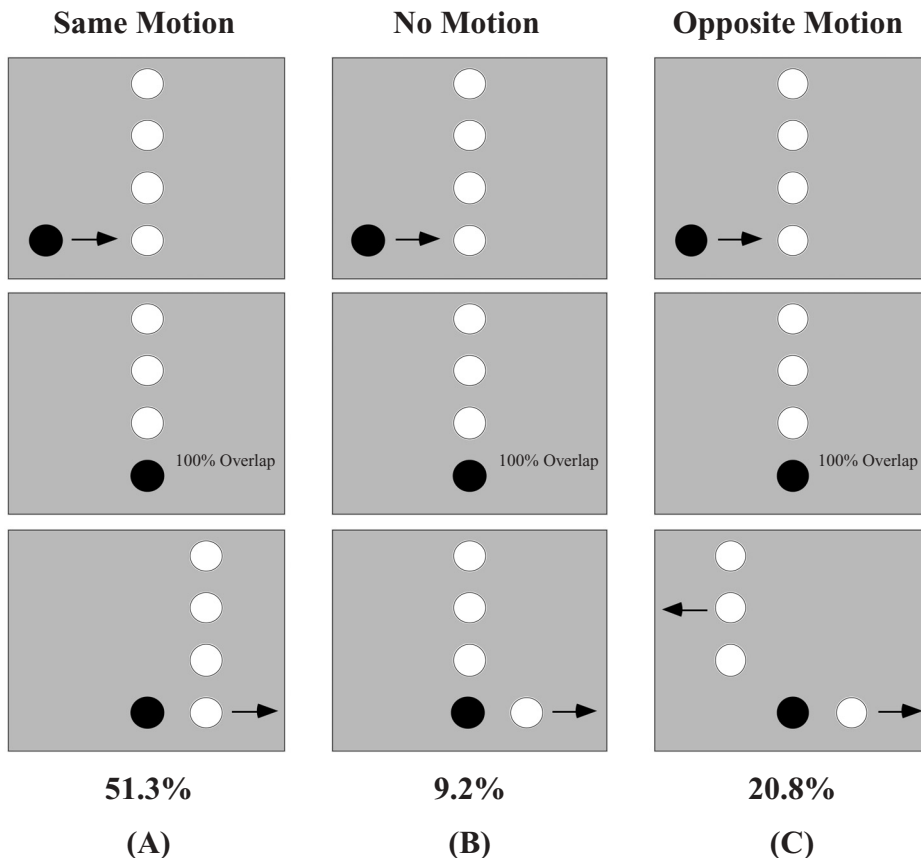


Figure 6. Depictions of the three conditions used in Experiment 2A and the resulting percentages of perceived causal launching. See the text for details. The result of the baseline condition is presented in Figure 7.

the context disks simply remained stationary in the center of the screen throughout the trial. We hypothesized that even these simple noninteracting (and certainly noncausal) context objects could influence the perception of causality. In particular, we predicted that the column of disks would initially group with the central object (B) in the test event and that B would then be seen to move (or not move) along with the context disks. Thus, in the no-motion condition, B would also be perceived as stationary, forcing A to be perceived as passing right through it, in continuous motion (noncausal passing). In the same-motion condition, B would be perceived as moving along with the moving context disks, forcing A to be perceived as stopping in the center of the display just as B starts moving (causal launching).

An *opposite-motion condition* (Figure 6C) was also included, in which the additional context disks remained stationary and aligned with B until the moment of full overlap in the test event, at which point the context items began moving in the opposite direction from B (i.e., back toward the initial horizontal location of A). This condition allowed us to separate any effects of the initial phase of motion (when the context items were grouped with B in the center of the display) from the final phase of motion (when the context disks were explicitly ungrouped from both A and B by moving separately).

Method

The same 12 observers as those from Experiment 1 participated in a separate session, with order fully counterbalanced. Each trial involved two components: the ambiguous full-overlap test event and three additional context disks. The test event was identical to that in Experiment 1 and was presented 8.94° above the lower display border. The three context disks were the same size and color as the initially stationary disk (B) in the test event and were initially positioned in a column directly above B, to form a column of four identical, regularly spaced disks, each separated from its neighbor(s) by 1.38° of blank space between their nearest edges. The context disks remained stationary in this position until the moment of complete overlap in the test event.

In the no-motion condition, the context disks remained stationary in the center of the screen throughout the trial. In the same-motion condition, the context disks then started moving (always as a column) in the same direction and at the same speed as B, eventually stopping when B stopped, near the right display border. In the opposite-motion condition, the context disks always moved when B moved but did so in the opposite direction (although again at the same speed) in such a way that they eventually stopped near the leftmost edge of the display, at the same horizontal position as that at which A had initially been presented. In the baseline condition, the test event was presented in isolation in the center of the display (11.87° from the upper display border). The observers completed 20 trials in each of the four conditions, for a total of 80 trials; all trials were presented in a random order, unblocked.

Results

The percentage of trials in each condition that were perceived as causal launching are shown at the bottom of each column in Figure 6. A single-factor repeated measures ANOVA revealed a significant effect of condition [$F(3,33) = 19.75, p < .001$]. Planned comparisons re-

vealed that the common motion of context disks strengthened the causal perception of the test event. The observers perceived causal launching more than five times as often in the same-motion condition (51.3%) as in both the baseline condition [6.7%; $t(11) = 6.16, p < .001$] and the no-motion condition [9.2%; $t(11) = 5.50, p < .001$], which didn't themselves differ [$t(11) = 0.86, p = .410$]. The opposite-motion condition yielded an intermediate degree of causal perception (20.8%), which was smaller than that in the same-motion condition [$t(11) = 3.72, p = .003$] and marginally greater than that in the baseline condition [$t(11) = 2.15, p = .055$].

Discussion

These results reveal a new type of contextual grouping effect on causal perception. Whereas our previous results demonstrated an influence of unambiguously causal contextual events (causal capture), this experiment demonstrated a simpler and more direct form of contextual grouping, wherein the context objects themselves were not perceived as causal, yet could still promote the perception of causality in the ambiguous test event. Why would these individual objects influence the perception of causality in this manner? It seems clear that in our displays, the context objects were grouped with the initially central item in the test event via good continuation and proximity (in the initial configuration) and by common motion (in the same-motion condition). Thus, these results demonstrate for the first time that the perception of causality can operate over perceptual *groups*, in addition to discrete objects.

The intermediate results of the opposite-motion condition begin to reveal some of the limits of this grouping process. Because this condition resulted in less than half the degree of causal perception as that observed in the same-motion condition, we conclude that the common motion per se—and not just the existence of any motion—greatly contributes to the grouping effect. On the other hand, the existence of a modest increase in causal perception even in the opposite-motion condition reveals that the mere existence of correlated motion can also promote causality. (Note that 20.8% of causal launching perceived here is roughly comparable with the 25.8% launching observed in the unconnected condition in Experiment 1, which also involved opposite motions in the test event and the context objects.) Further research will be required to determine why opposite-motion contexts are sufficient by themselves to induce a modest amount of causal perception in the test event, but we can speculate that the two directions of motion are still grouped to some degree because (1) the motions are still perfectly temporally synchronized and (2) the motions still occur in the same horizontal plane and emanate from a central point.

EXPERIMENT 2B

Varying the Number of Context Disks

In Experiment 2A, we hoped to maximize the likelihood of perceptual grouping by including a full column of

context disks. The strength of the resulting grouping effect, however, led us to question whether a similar effect might be observed with even a single contextual object that is grouped via common motion with the items in the test event. To explore this question, we replicated the same-motion condition of Experiment 2A but varied the number of context disks from one to three.

Method

The same 12 observers as those in Experiment 2A participated. Each trial was identical to those in the same-motion condition in Experiment 2A, except that the number of context disks was varied. The *3-disk condition* was identical to the same-motion condition in Experiment 2A. The *2-disk condition* was identical except that the uppermost context disk was never presented (leaving an initial column of three disks, including the initially stationary Disk B from the test event). The *1-disk condition* was identical except that only the lowest context disk was presented (leaving a column of two disks, including B). The observers completed 20 trials in each of the three conditions, for a total of 60 trials; all the trials were presented in a random order, unblocked.

Results and Discussion

The percentage of trials in each condition that were perceived as causal launching are shown in Figure 7. (The no-context baseline condition from Experiment 2A is also shown here and is used in the analyses below, since the two experiments employed the same observers.) As is clear from the graph, the presence of context objects promoted the perception of causality, but the number of context objects had no effect [$F(2,22) = 0.44, p = .652$]. A planned comparison between the 1-disk condition and the

no-context baseline condition revealed that a significant grouping effect occurred even in the 1-disk condition and gave rise to significantly more causal perception [$t(11) = 5.99, p < .001$]. Thus, the inclusion of just a single context object can influence the perception of causality in an ambiguous event.

EXPERIMENT 2C Effects of Proximity

The results of Experiment 2B suggest that even a moderate amount of grouping from a single contextual object can greatly influence causal perception. This effect seems likely to have been due to the common motion, as was discussed above, but may also have been strengthened here by the close proximity of the context object to the initially stationary object in the test event. To test whether proximity was indeed driving this grouping effect, we replicated the 1-disk condition from Experiment 2B but varied the distance between the single context disk and the test event.

Method

The same 12 observers as those in Experiment 2A participated. The trials were identical to those in the 1-disk condition of Experiment 2B, except that the vertical distance between the test event and the context disk was varied in three conditions. The *near condition* was identical to the 1-disk condition in Experiment 2B. In the *far condition*, the nearest edges of the context disk and the initially stationary Disk B in the test event were 7.70° apart. In the *intermediate condition*, the nearest edges of the context disk and Disk B in the test event were 4.47° apart. The observers completed 20 trials in each of

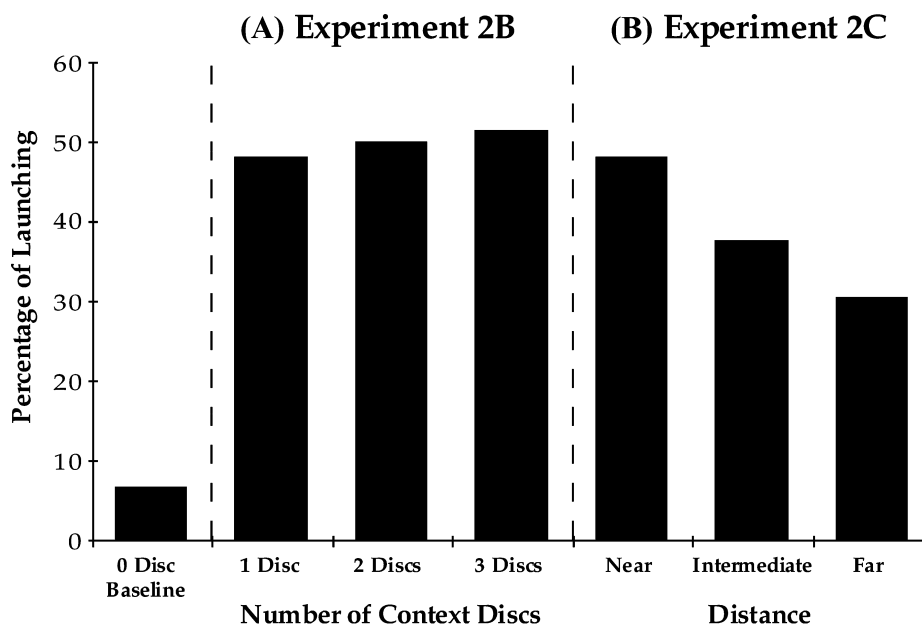


Figure 7. The percentage of perceived causal launches in each condition of (A) Experiment 2B and (B) Experiment 2C. “Baseline” refers to the Baseline condition in Experiment 2A, wherein the full-overlap event was presented in isolation. Because the same observers participated, we used this result for our analysis.

the three conditions, for a total of 60 trials; all the trials were presented in a random order, unblocked.

Results and Discussion

The percentage of trials in each condition that were perceived as causal launching are shown in Figure 7. In contrast to the null effect of the number of context disks in Experiment 2B, the graph suggests that the proximity of a single disk did influence the degree of causal perception in the ambiguous test event [$F(2,22) = 11.43, p < .001$]. Further planned comparisons indicated that causal perception was more likely with increasing proximity: the near condition (47.9%) gave rise to more causal perception than did the intermediate condition [37.5%; $t(11) = 3.23, p = .008$], which in turn gave rise to marginally more causal perception than did the far condition [30.4%; $t(11) = 2.16, p = .054$]. Note, however, that a significant effect of grouping by common motion occurred even in the far condition, which still gave rise to more causal perception than did the no-context baseline condition in Experiment 2A [$t(11) = 4.45, p = .001$].

The significant effect of even a single relatively distant contextual object in this experiment is a testament to the power of grouping by common motion in causal perception. At the same time, the results of this experiment demonstrate that grouping by proximity also plays an important role. Thus, at some level, the underlying units of causal perception can clearly be perceptual groups. But how do such grouping effects arise? The remaining experiments pursued the idea that these effects are realized by the automatic spread of attention within a group.

EXPERIMENT 3 Reverse Causal Capture

All of our previous studies of causal perception have explored the ways in which contextual information can promote the perception of causality in an ambiguous full-overlap test event that is often perceived in isolation as a noncausal pass. But what about the reverse situation?

Could contextual information also *attenuate* the perception of causality in a test event that is typically seen in isolation as a causal launch? None of the previous studies has addressed this question (although see the *camouflage* studies of Michotte 1946/1963, as described in the General Discussion section). Here, we demonstrate for the first time that *reverse causal capture* of this type is possible.³

In fact, we still think it unlikely that such contextual information could attenuate the perception of causality in a completely unambiguous, nonoverlapping launch event. However, the present experiment was made possible by weakening the launching event so that the two disks only *partially* overlapped during the event (by 67% of their diameters). This intermediate condition is still perceived in isolation as causal launching on an overwhelming majority of trials (Scholl & Nakayama, 2004), yet the partial overlap now still allows contextual information to attenuate this robust causal perception.

We explored this phenomenon using four conditions. In the baseline *partial-overlap launching* condition, the 67% overlap event was presented in isolation (Figure 8A). In our three other conditions, a second context event was added to the display. Whereas the context events in our previous experiments were always unambiguous nonoverlapping causal launches, the context events employed here were now *noncausal* full-overlap events. Since we had found in Experiment 2C that the distance between events can greatly influence the impact of contextual information on causal perception, this context event was included at three separate distances from the partial-overlap test event. As in Experiment 2C, these constituted our *near context*, *intermediate context*, and *far context* conditions (the first of which is depicted in Figure 8B).

Method

Ten observers participated; none had participated in any of the previous experiments. In the partial-overlap launching condition, the single 67%-overlap event was presented in isolation, roughly in the center of the display (so that the highest point of each disk was 10.18° from the upper display border). The initially stationary disk (B) in this test event was always bright green, whereas the initial moving




Figure 8. Depictions of (A) the partial-overlap launching condition and (B) the near context condition in Experiment 3. In the partial-overlap event, two disks overlapped 67%. In the near context condition, the observers had to judge the lower event, wherein the launching disk was blue.

disk (A) was always bright blue. (In this experiment and the following experiment, the observers were thus simply instructed to always report their percept of the event that contained the blue disk.) In each of the other conditions, this event was presented along with a full-overlap event (as described in the General Method section). These two events were always temporally synchronized, so that the moment of maximal (67%) overlap at the center of the display in the partial-overlap test event occurred at the same moment as that at which the two disks in the context event were fully overlapped in the center of the display. The lowest point of each disk in the test event was always 8.94° from the lower display border, and the full-overlap context event always appeared above the test event. The nearest edges of the disks in the two events were separated by 1.38° in the near context condition, 4.47° in the intermediate context condition, and 7.70° in the far context condition. The observers completed 20 trials in each of the four conditions, for a total of 80 trials; all the trials were presented in a random order, unblocked.

Results

The percentage of trials in each condition that were perceived as causal launching are shown in Figure 9. We first note that the baseline partial-overlap launching condition is indeed aptly named: The observers perceived causal launching on a majority of trials (87%) when this event was presented in isolation. This robust causal perception was weakened, however, by the presence of the additional full-overlap context event. This weakened causal perception was most evident in the near context condition, when the two events were closest: Here, the degree of causal launching was reduced by 30% [from 87% to 57%; $t(9) = 3.97, p = .003$]. Moreover, causal launching was reduced by more than 15% even in the weakest Far Context condition [from 87% to 71.5%; $t(9) = 2.99, p = .015$]. A single-factor repeated measures ANOVA confirmed that the varying degrees of proximity between the events influenced causal perception [$F(2,18) = 5.47, p = .014$], and further planned comparisons revealed a significant difference between the far and the intermediate conditions [$t(9) = 2.40,$

$p = .040$] and a marginal difference between the intermediate and the near conditions [$t(9) = 2.14, p = .061$].

Discussion

These results demonstrate that noncausal contextual information can actually attenuate the perception of causality in an event that is perceived as robustly causal in isolation. Since this effect is in many ways the opposite phenomenon to that studied by Scholl and Nayakama (2002), we dub this the *reverse causal capture* effect. Because of the common motion, temporal synchrony, and proximity between the two events, it again seems natural to think of this effect as resulting from a type of perceptual grouping between the two events. In this sense, this experiment emphasizes the power of grouping on the perception of causality: In reverse causal capture, this grouping is strong enough to override even the normally robust perception of causal launching. In the following experiment, this reverse causal capture effect was employed in a direct demonstration of the role of attention in causal perception.

EXPERIMENT 4A

Effects of Attention on Perceived Causality

The previous experiments have demonstrated in various ways that perceptual grouping can greatly impact the perception of causality, and we have hinted that such effects may be driven by the allocation of attention in dynamic visual scenes. In particular, we hypothesize that attention automatically spreads within image elements defined by grouping cues (cf. Scholl, 2001) and that this may be *why* we see the perceptual effects of grouping—that is, that certain disks *go together* (Driver, Baylis, Russell, Turrato, & Freeman, 2001). In this experiment, we directly demonstrate an effect of attention on causal perception, by

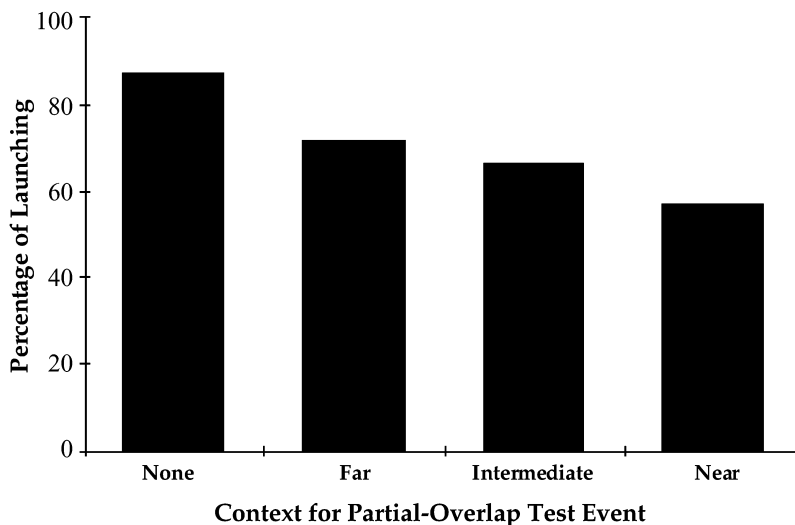


Figure 9. The percentage of perceived causal launching in each condition in Experiment 3.

manipulating the focus of attention in various conditions within identical displays.

In many ways, it was Experiment 3 in this study that resulted in the strongest grouping effect, since there it was found that grouping cues could attenuate causal perception even in an event that is seen as robustly causal in isolation. Accordingly, the present experiment involved an adaptation of this reverse causal capture display. In particular, the far context condition was used, wherein the items were widely separated, leaving room for various loci of attention between the two events. As in that condition, each trial involved two relatively widely separated events: One was the partial-overlap launching test event, and the other was the full-overlap noncausal context event. In this experiment, either of these events could appear above the other, and the observer's task was simply to always report the perceived causal status of the event with the blue disk—which was, in fact, always the partial-overlap test event (presented here as either the upper or the lower event). The measure of interest in this experiment, as in Experiment 3, is the degree of *attenuation* of perceived causality in the normally causal test event by the noncausal context event. In Experiment 3, this attenuation caused the test event to be seen as causal 15% less often than when it was presented in isolation.

Unlike Experiment 3, which was conducted under free-viewing conditions, we now manipulated, in three conditions, where the observers attended while observing the display. In the *attend-test* condition, observers attended directly to the test event. In the *attend-context* condition, the observers attended directly to the context event. And in the *attend-between* condition, the observers attended to the empty space directly between the two events. In all cases, the observers reported the perceived causal status of the test event (with the blue disk), regardless of where they were attending. The different attention conditions were all randomized within this experiment and were manipulated as follows. Before each trial, an arrow appeared toward one side of the display and flashed several times. The horizontal path pointed to by the arrow indicated the *row* of the display to which the observers should attend during that trial. We did not explicitly manipulate eye movements but told the observers that they were free to move their eyes but should do so only in the horizontal path of the display indicated by the initial arrow.

Method

The same 10 observers as those in Experiment 3 participated in a separate session, with order fully counterbalanced. The basic display was identical to that in the far context condition of Experiment 3, with one exception: The test event was positioned below the context event only in half of the trials, whereas in the other half the test event was positioned above the context event. The positions of the events never varied: The highest point of the upper event was always 5.46° from the upper display border, and the lowest point of the lower event was always 8.94° from the lower display border.

Before each trial began, a bright yellow arrow (presented as a rotated triangle 'pointing' to the right, inscribed in a rectangle subtending 0.73° × 0.99°) flashed for 1,300 msec in the near left edge of the screen, 0.69° from the left border of the screen. The vertical

position of the arrow was adjusted so that it occurred (1) at the same vertical position as the test event (essentially pointing at the horizontal path through which the test event objects would move), (2) the same vertical position as the context event, or (3) equally spaced between these locations (essentially pointing at a horizontal path centered between the two events). The observers were instructed to attend only along the horizontal path pointed to by the arrow, regardless of whether this path contained an event or was centered between the two events. In all cases, both events were still clearly visible, and the observers reported the perceived causal status of the event that contained the blue disk (i.e., the partial-overlap test event). The observers completed 40 trials in each of the three conditions, for a total of 120 trials. The test event appeared above the context event on half of the trials in each condition and below the context event in the other half. (Thus, the observers were attending to the upper row on half of the attend-test trials and to the lower row on the other half of the attend-test trials, and similarly for the attend-context trials.) All the trials were presented in a random order, unblocked.

Results

The percentages of trials in each condition that were perceived as causal launching are presented in Figure 10A. As is clear from this graph, the locus of attention had a considerable impact on the perception of causality in the test event [$F(2,18) = 40.73, p < .001$]. Essentially, the perceived causal status of the test event was dominated by the causal status of the event in the horizontal row to which the observers attended. In isolation, the partial-overlap test event was seen as causal launching on an overwhelming majority of trials (87% in Experiment 3); accordingly, the observers were likely to perceive causal launching in this test event (on 74.8% of trials) when attending directly to it. In contrast, when the observers attended to the noncausal full-overlap context event, this same test event was perceived as causal launching on only 20.5% of the trials. An intermediate degree of causal launching was perceived when the observers attended between the two events (30.5%). Additional planned comparisons revealed a significant difference between the degrees of perceived causal launching in the attend-test and the attend-between conditions [$t(9) = 6.16, p < .001$] and a marginal difference between the attend-between and the attend-context conditions [$t(9) = 2.25, p = .051$].⁴

Further analyses explored the impact of the relative positions of the context and the test events (i.e., which was below the other). In general, the reverse causal capture effect (i.e., the degree of attenuation of perceived causal launching) was greater when the context event was presented in the lower position (36.5% perceived causal launching) than when the test event was presented in the lower position [48%; $t(9) = 3.50, p = .007$]. Furthermore, this positional effect occurred in both the attend-context condition [14.5% vs. 26.5%; $t(9) = 2.48, p = .035$] and the attend-between condition [22.5% vs. 39.5%; $t(9) = 3.13, p = .012$] but did not occur in the attend-test condition [72.5% vs. 78%; $t(9) = 1.30, p = .227$].

Discussion

The results of this experiment confirm that the locus of attention can have a strong effect on the perception of

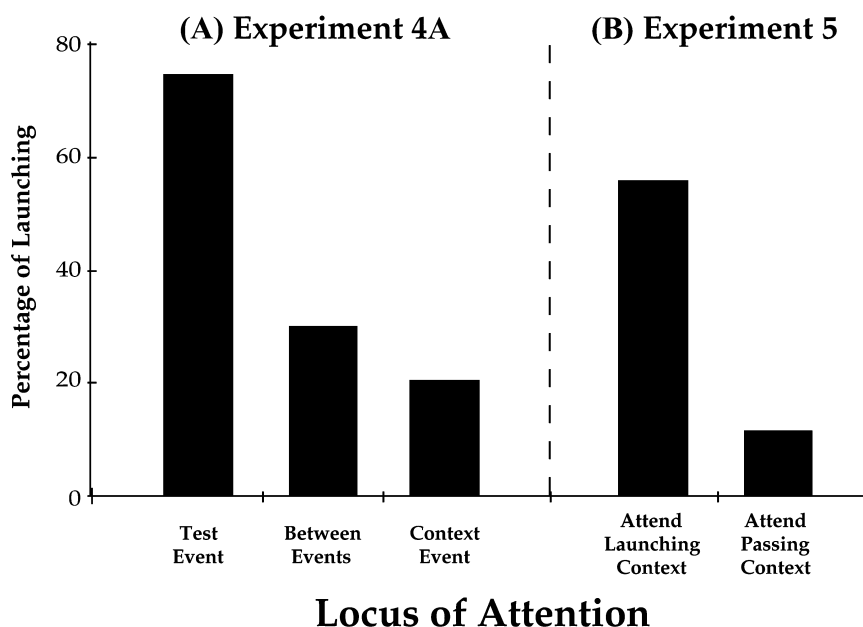


Figure 10. The percentage of perceived causal launching in each condition of (A) Experiment 4A and (B) Experiment 5.

causality. Here, the degree of perceived causal launching *in the very same display* was modulated by more than 50% on the basis only of where the participants attended. (We also note, in passing, that this effect is perceptually salient, as can perhaps be experienced in the on-line demonstrations.) This result thus demonstrates that the *input* to the processes that mediate causal perception are filtered through attention, and the implications of this will be discussed in the General Discussion section.

The results also revealed an interesting anisotropy between the upper and the lower visual fields, which is also reminiscent of attentional effects. Whenever the full-overlap context event was located below the test event, it had a greater influence, so that the resulting reverse causal capture effect was greater. Conversely, when the test event was located below the context event, reverse causal capture was weaker. This pattern was clearest in the attend-between condition, wherein the two events were always equidistant from the locus of attention. Here, the context event had a greater effect when it was in the lower visual field than when it was in the upper visual field. This lower-field advantage has been found in several other studies of attentional resolution, involving feature discrimination (He, Cavanaugh, & Intriligator, 1997), multiple-object tracking (He, Cavanaugh, & Intriligator, 1996), and change detection (Intriligator, He, & Barton, 1998). Thus, we see not only a direct main effect of attention on the perception of causality, but also more subtle effects due to anisotropies in the resolution of attention.

EXPERIMENT 4B

Unconfounding Attention and Eccentricity

Although the manipulations in Experiment 4A surely did involve attention, they also perfectly confounded at-

tention and eccentricity. Indeed, our manipulation of attention was implemented directly by cuing the observers to keep their eyes fixated on different events (or on the empty space between events) while continuing to judge a test event. Because of this, fixating a context event (while continuing to judge a test event) always involved making the test event more retinally eccentric. Could our alleged effects of attention be fully explained by such retinal eccentricity? If so, perhaps our context events were entirely unnecessary; merely moving the partial-overlap test event into the relative periphery might decrease causal perception, regardless of whether this resulted in fixating a context event or fixating nothing at all. One reason for pursuing this alternate explanation is that such eccentricity effects surely do exist. Michotte (1946/1963, Experiment 7), for example, demonstrated effects of fixation position with single events, and these can be readily appreciated in our displays simply by viewing the Web-based demonstrations in the extreme periphery.

Here, we directly unconfound attention to context events and retinal eccentricity by varying whether or not a context event was present. As in Experiment 4A, the observers always reported the perceived causal status of a partial-overlap test event, while either fixating that test event or fixating a more peripheral location, which could contain either a full-overlap context event or simply a fixation point. If the results of Experiment 4A did, in fact, reflect the operation of attention, we would expect the presence of the context event to matter when eccentricity and fixation are held constant.

Method

This experiment was identical to Experiment 4A, except as noted here. Ten observers participated; none had participated in any of the

previous experiments. There were two display conditions. The *context event* display was identical to that in Experiment 4A. The *context fixation* display was identical, except that a single static yellow fixation cross (subtending 1.2°, horizontally centered in the display) replaced the context event. As in Experiment 4A, the relative location (upper vs. lower position) of the test event and the context stimulus were fully counterbalanced in these displays across trials. There were two attention conditions: An attentional cue (implemented as in Experiment 4A) instructed the observers to fixate either the test event or the context stimulus. The observers completed 20 trials in each of the four conditions (two attention conditions crossed with two display conditions), in a fully randomized order, unblocked.

Results and Discussion

When the context stimulus was simply a fixation point, there was no effect of the attention manipulation and no reverse causal capture: The observers tended to perceive causal launching in the test event on a large majority of trials regardless of fixation position [91.5% when fixating the test event vs. 87% when fixating the context fixation point; $t(9) = 1.22, p = .253$]. In contrast, when the very same fixation positions were used in the presence of the full-overlap context event, we observed a robust effect of attention and considerable reverse causal capture, replicating Experiment 4A [87.5% perceived launching when fixating the test event vs. 53% when fixating the context event; $t(9) = 4.92, p = .001$]. When fixation was held constant, so that observers fixated the context position and made the test event relatively retinally eccentric, we observed a large significant effect of the presence of a context event at fixation [87% perceived launching with fixation vs. 53% with the context event; $t(9) = 4.43, p = .002$]. These results support the idea that the presence of the context event in Experiment 4A was crucial and that these results truly reflect the operation of attention, and not simply retinal eccentricity.

EXPERIMENT 5 Effects of Attention to Competing Context Events

In Experiments 4A and 4B, we demonstrated effects of attention by manipulating whether observers fixated the test event or not, and they often had to report the perceived causal status of an event when it was not being directly fixated. Here, we present a complementary demonstration of the influence of attention on causal perception wherein observers are always fixating a centrally located test event but must also attend to only one of two flanking context events that are in competition. The test event in this experiment was always the ambiguous full-overlap event (Figure 2). The *explicit launching* context event had no overlap, so that it was likely to induce causal capture in the test event. The *explicit passing* context event, in contrast, involved a full-overlap event with no color change that when attended, is overwhelmingly seen as passing. Both context events were present on all trials, flanking the test event, which was always fixated, but the observers were instructed on each trial to attend to only one of the context events. (To our knowledge, this is the first study of causal

perception in which multiple contextual sources of information are in competition.) This allows us to test for effects of attention when both the display and the fixation position are held constant.

Method

The same 10 observers as those in Experiment 4B participated in a separate session, with order fully counterbalanced. Each display involved the same three events—a test event and two context events—each involving an initially central green disk and a red disk that moved from left to center. The test event was a full-overlap event, now vertically centered in the display. One context event was an unambiguous launching event. The second context event was an *explicit passing* event that differed from the full-overlap event in two respects: (1) The stationary disk was always green, and the moving disk was always red, and (2) during the period of intersection, the stationary green disk was always drawn on top of the moving red disk. This yielded a percept of a moving red disk unambiguously passing behind a stationary green disk. The context events flanked the test event (one above and one below, fully counterbalanced), so that the nearest edges of each adjacent pair of events were always separated by 1.92°. The observers were instructed to always report the perceived causal status of the central test event, but to attend on each trial to one of the two context events, which blinked momentarily before the motion began. The observers completed 20 trials, attending to each of the two context events (10 in the upper position and 10 in the lower position), for a total 40 trials, presented in a fully randomized order, unblocked.

Results and Discussion

In this experiment, the fixation position was constant, and the display was identical across trials, except for the relative positions of the two context events. Nevertheless, attention greatly influenced causal perception: As is depicted in Figure 10B, the very same central fixated full-overlap test event was perceived as causal launching much more often when the observers attended to the explicit launching context event than when they attended to the explicit passing context event, even though both were always present in the display and were equidistant from fixation [11% vs. 55.5%; $t(9) = 5.93, p < .001$]. These results provide further evidence that attention can directly influence the perception of causality even when the judged event is directly fixated: Attention can determine not only *whether* a context event will influence causal perception in isolation, but also *which* context events will affect perceived causality in what is perhaps a more ecologically valid situation, with multiple events in the same scene.

GENERAL DISCUSSION

Michotte's classic experiments on the perception of causality were primarily an attempt to determine the specific stimulus factors that the visual system uses to detect causal relations. The experiments reported in this article shared this general goal but explored a wider array of contextual image cues, focusing on the roles of perceptual grouping and attention.

Experiment 1 demonstrated a role for perhaps the simplest form of perceptual grouping: actual physical connectedness between objects. Experiments 2A, 2B, and 2C then generalized this influence of grouping to other cues,

including proximity, good continuation, and common motion. Here, we found evidence that even a single additional object that moved along with one of the objects in the ambiguous test display could influence the perception of causality. Experiment 3 introduced a new type of context effect, reverse causal capture, wherein contextual information could *attenuate* the perception of causality in a stimulus typically seen in isolation as causal launching. Throughout these experiments, we hinted that the observed effects of perceptual grouping might, in fact, be driven by the differential allocation of visual attention.

In Experiments 4 and 5, we directly demonstrated such a role for attention. In Experiment 4, we again employed the reverse causal capture display, but now showed that the degree of perceived causal launching in the very same display could be modulated by over 50% on the basis of where the observers attended. Similarly, we showed in Experiment 5 that attention could influence perceived causality even when the judged event was always fixated, when two context events were in competition.

These effects of grouping and attention on the perception of causality are important for several reasons.

The Role of Context in Causal Perception

Collectively, these results demonstrate that the perception of causality is sensitive not just to local information within individual events, but that it can also take a wide variety of contextual information into account, driven by attention and perceptual grouping. This type of context sensitivity is a hallmark of many types of visual processing, both low level and high level. At the level of individual neurons, physiologists have shown that the responsiveness of a neuron to a stimulus in its receptive field can be both facilitated and inhibited by information from neighboring stimuli beyond the receptive field (e.g., Allman, Miezin, & McGuinness, 1985; Levitt & Lund, 1997; Nelson & Frost, 1978). Similar effects are observed in many aspects of conscious perception, wherein objects and events from different parts of a visual scene can *entrain* each other (e.g., in apparent motion; Ramachandran & Anstis, 1983).

The causal capture effect employed here is, of course, an example of such a context effect on the perception of causality, but the present results go beyond causal capture in several ways. Scholl and Nayakama (2002) suggested that causal capture was intrinsically an effect of capture *by* causality—that is, that the capture effect was driven by additional causal *events*. The present results, in contrast, suggest that this view is overly constrained in two important ways. First, Experiment 3 demonstrated that information even from *noncausal* events (e.g., full-overlap passing) can influence the perception of causality, via reverse causal capture. Second, Experiment 2B demonstrated that the contextual information taken into account by the visual system when inferring a causal interaction need not occur in the context of an *event* at all: Even a single isolated context *object* can have an overwhelming effect on the perception of causality, if it moves in a way that groups it with other objects in the judged event.

Such context effects on the perception of causality actually have their roots in some of Michotte's (1946/1963) earliest work. Among the more than 100 experiments reported in his seminal book are several demonstrations of ways in which additional objects can influence the perception of causality.⁶ In particular, Michotte described several demonstrations wherein additional contextual information could attenuate the perception of causality in a simple launching event. The essence of these manipulations is to provide a competing “explanation” for the motion of the putatively “launched” second disk (B). For example, in one of his experiments (Experiment 21), Disk A approached and launched Disk B in the normal fashion. However, this typical motion pattern was preceded by a series of to-and-fro movements by Disk B between its ordinary stopping place (near the right border of the display) and its initial position near the center of the display. This series of movements was synchronized with the main launching event, so that critical movement of Disk B (i.e., upon A's arrival) was simply an identical continuation of the motions in which B had already been engaged. In this situation, participants tended not to perceive a causal relationship between A and B. (For another conceptually similar experiment, see Michotte, 1946/1963, Experiment 20.)

This type of *camouflage* effect on the perception of causality is clearly similar to the context effects reported here: Like the reverse causal capture experiment (Experiment 3), it demonstrates that contextual information can attenuate the perception of causality. However, the contextual information in Michotte's experiment was inherently *temporal* in nature—an object's earlier behavior influences its perception at a later time—whereas the grouping effects reported above in Experiment 2 primarily involve *spatial* contextual interactions, wherein an object is influenced by the behavior of a spatially distinct object at the same time.

What Are the Units of Causal Perception?

What do the context effects described above tell us about the nature of the perception of causality? In essence, such results may require changes in what we think of as the underlying *units* of causal perception. A critical issue for any perceptual (or cognitive) process is the nature of the units over which that process can operate. Ever since Michotte, we have simply assumed that the units of perceived causality were individual objects—as is true of most collisions in the real world (including Hume's billiard balls). Indeed, to our knowledge, nobody has previously explored the units of causal perception experimentally. This is true even though Michotte was in some sense a Gestalt psychologist and proposed a Gestalt theory of causal perception.

The present experiments show that other types of units are also possible, however—in particular, perceptual groups. This is perhaps most clearly demonstrated in Experiment 2A, wherein the behavior of an additional column of disks largely determined how the ambiguous full-overlap

test event (involving A and B) would be perceived. When the additional disks (which we will collectively call C) remained stationary at the center of the screen throughout the trial, the observers tended to see noncausal passing. When C always stayed aligned with B (the launched disk) in the main event, the observers tended to perceive causal launching. This seems best explained in terms of grouping: B and C are always grouped together throughout the event—by proximity, similarity, good continuation, and common motion. As a result, the initially central disk is perceived to do whatever the C disks do: When they remain stationary, it does too (resulting in the passing percept); when they move upon A's arrival at the center of the display, it does too (resulting in the launching percept). In essence, when C moves upon A's arrival, the primary event is *not* an ambiguous event involving A and B, but an unambiguous launching event involving A and the *group composed of B and C*. (Since A does not fully overlap C, there is no ambiguity to cause a passing percept.) In this sense, the underlying units of causal perception can be groups, in addition to individual objects. This increase in the scope of the potential units of causal perception mirrors similar shifts that have occurred in the units of "object-based" attention (e.g., Driver & Baylis, 1998; Scholl, 2001, section 3.3) and of perceived animacy (Bloom & Veres, 1999).

Attention and Causality

We suggest that all of the grouping effects reported in this article may be explained in terms of the operation of visual attention. This was clearest in Experiments 4 and 5, in which the degree of perceived causal launching in the very same display was modulated by more than 50% only on the basis of where participants attended. However, we can also interpret the other grouping effects in terms of attention. In doing so, we follow other recent researchers who have raised the possibility of identifying grouping processes with mechanisms of attention. Although the classical studies of perceptual grouping never referred directly to attention, a long tradition in psychology has considered that these two factors are closely related (e.g., Neisser, 1967). Often the direction of this relationship is unclear: Some researchers have argued that the spread of attention can be explained by the operation of grouping processes, whereas others have argued that groups are perceived precisely because of the operation of attention (e.g., Driver et al., 2001; Scholl, 2001). For example, one recent review suggests that

the introspection that a subset of dots in a Gestalt display "belongs together" may arise precisely because when trying to pick out one of these dots, you tend to pick out also those dots which are grouped with it. In other words, they may seem to "belong together" precisely because you tend to attend them together. (Driver et al., 2001, p. 65)

Thus, there may simply be a single set of processes that are responsible for both the standard object-based attention effects and the phenomenology of perceptual grouping.

This perspective helps to unify the experiments presented in this article: Experiments 4 and 5 demonstrated effects on causal perception of *voluntary* attentional allocation (via the task instructions and the cued row), whereas the other experiments demonstrated effects of *automatic* attentional spread, induced by connectedness (in Experiment 1) and other grouping cues (in Experiment 2). Thus, the seemingly disparate experiments reported here may in fact all reflect the same underlying process: the voluntary or automatic spread of attention.⁷

The fact that causal perception can be affected by the allocation of attention may also have other theoretical implications, beyond the relation to grouping processes. Following Michotte, some contemporary researchers have suggested that processes that mediate the perception of causality are part of an encapsulated (and possibly innate) *module* in the visual system (e.g., Leslie, 1986; Leslie & Keeble, 1987; Scholl & Tremoulet, 2000). This view suggests that causal perception should be a largely data-driven system, determined primarily by the incoming visual input, and not by the observer's intentions or beliefs. But what is the nature of this input? Some early discussions seem to suggest that the modular processes that mediate causal perception may operate on very early visual representations, perhaps in parallel across the visual field. Here, we have shown that this is not the case: The inputs to such a module are at least filtered through attention. Of course, this result does not necessarily impact the degree to which causal perception is a *lower level* process: After all, attention itself can affect even some of the earliest levels of visual processing (Kastner & Ungerleider, 2000). So, can beliefs and intentions affect the visual processes that are responsible for inferring the existence of a mechanical causal relation? Our results at least demonstrate that although beliefs and intentions may not be able to directly affect the nature of such processing itself, they can affect the *input* to such processes, via the information gated through the voluntary allocation of attention. (In this way, effects of attention do *not* imply that the processes responsible for causal perception are *cognitively penetrable*; cf. Pylyshyn, 1999).

Conclusions

Despite the longstanding interest by psychologists in the perception of causality, relatively few recent studies have explored the principles which underlie it. We suppose that this may be for two primary reasons. First, studies of causal perception have always tended to be curiously disconnected from more mainstream perception research—as if the perception of causality was somehow an epiphenomenon. Second, we think that many researchers have suspected that Michotte himself characterized most of the fundamental *rules* of causal perception, in over 100 experiments reported in his landmark book (1946/1963), and that more recent work can delve only into more specialized cases.

The experiments reported here belie both of these reasons, demonstrating that causal perception interacts richly

with other visual processes, such as grouping and attention. Moreover, effects of grouping and attention represent entirely new classes of principles, which go well beyond the factors studied by Michotte. We thus propose that the rules that underlie the perception of causality are both more complex and more interesting than has been previously suspected and that continued investigation of these rules will help us understand not only how the visual system infers the existence of a causal relation, but also how such processes are deeply connected to many other aspects of visual perception.

REFERENCES

- ALLMAN, J., MIEZIN, F., & MCGUINNESS, E. (1985). Direction- and velocity-specific responses from beyond the classical receptive field in the middle temporal visual area (MT). *Perception*, **14**, 105-126.
- BLOOM, P., & VERES, C. (1999). The perceived intentionality of groups. *Cognition*, **71**, B1-B9.
- BOYLE, D. (1960). A contribution to the study of phenomenal causation. *Quarterly Journal of Experimental Psychology*, **12**, 171-179.
- DARBY, D. (2002). *MacStim* (Version 2.6.8) [Computer software]. West Melbourne, Australia: Author.
- DRIVER, J., & BAYLIS, G. (1998). Attention and visual object segmentation. In R. Parasuraman (Ed.), *The attentive brain* (pp. 299-325). Cambridge, MA: MIT Press.
- DRIVER, J., BAYLIS, G., RUSSELL, C., TURATTO, M., & FREEMAN, E. (2001). Segmentation, attention, and phenomenal visual objects. *Cognition*, **80**, 61-95.
- GEMELLI, A., & CAPPELLINI, A. (1958). The influence of the subject's attitude in perception. *Acta Psychologica*, **14**, 12-23.
- GORDON, I., DAY, R., & STECHER, E. (1990). Perceived causality occurs with stroboscopic movement of one or both stimulus elements. *Perception*, **19**, 17-20.
- GUSKI, R., & TROJE, N. F. (2003). Audiovisual phenomenal causality. *Perception & Psychophysics*, **65**, 789-800.
- HE, S., CAVANAGH, P., & INTRILIGATOR, J. (1996). Attentional resolution and the locus of visual awareness. *Nature*, **383**, 334-337.
- HE, S., CAVANAGH, P., & INTRILIGATOR, J. (1997). Attentional resolution. *Trends in Cognitive Sciences*, **1**, 115-121.
- HUBBARD, T. L., BLESSUM, J. A., & RUPPEL, S. E. (2001). Representational momentum and Michotte's launching effect paradigm. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 294-301.
- HUME, D. (1960). *A treatise of human nature*. Oxford: Oxford University Press, Clarendon Press. (Original work published 1740)
- HUME, D. (1977). *An enquiry concerning human understanding*. Indianapolis: Hackett. (Original work published 1748)
- INTRILIGATOR, J., HE, S., & BARTON, J. (1998). Detection of object-change is more accurate in the lower visual field [Abstract]. *Investigative Ophthalmology & Visual Science*, **39**, S852.
- KASTNER S., & UNGERLEIDER, L. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, **23**, 315-341.
- KNILL, D., & RICHARDS, W. (1996). *Perception as Bayesian inference*. Cambridge: Cambridge University Press.
- KRUSCHKE, J. K., & FRAGASSI, M. M. (1996). The perception of causality: Feature binding in interacting objects. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 441-446). Mahwah, NJ: Erlbaum.
- LESLIE, A. M. (1986). Getting development off the ground: Modularity and the infant's perception of causality. In P. van Gest (Ed.), *Theory building in development* (pp. 405-437). Amsterdam: North-Holland.
- LESLIE, A. M., & KEEBLE, S. (1987). Do six-month-old infants perceive causality? *Cognition*, **25**, 265-288.
- LEVITT, J. B., & LUND, J. S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature*, **387**, 73-76.
- MARR, D. (1982). *Vision*. New York: Freeman.
- MICHOTTE, A. (1963). *The perception of causality* (T. Miles & E. Miles, Trans.). New York: Basic Books. (Original work published 1946)
- MICHOTTE, A., & THINÈS, G. (1991). Perceptual causality. In G. Thinès, A. Costall, & G. Butterworth (Eds. & Trans.), *Michotte's experimental phenomenology of perception* (pp. 66-87). Hillsdale, NJ: Erlbaum. (Original work published 1963)
- MORRIS, M. W., & PENG, K. (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality & Social Psychology*, **67**, 949-971.
- NATSOUKAS, T. (1961). Principles of momentum and kinetic energy in the perception of causality. *American Journal of Psychology*, **74**, 394-402.
- NEISSER, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.
- NELSON, J. I., & FROST, B. (1978). Orientation-selective inhibition from beyond the classic visual receptive field. *Brain Research*, **139**, 359-365.
- PALMER, S., & ROCK, I. (1994). On the nature and order of organizational processing: A reply to Peterson. *Psychonomic Bulletin & Review*, **1**, 515-519.
- PYLYSHYN, Z. W. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral & Brain Sciences*, **22**, 341-423.
- RAMACHANDRAN, V. S., & ANSTIS, S. M. (1983). Perceptual organization in moving pattern. *Nature*, **304**, 829-831.
- SCHLOTTMANN, A. (2000). Is perception of causality modular? *Trends in Cognitive Sciences*, **4**, 441-442.
- SCHLOTTMANN, A., ALLEN, D., LINDEROTH, C., & HESKET, S. (2002). Perceptual causality in children. *Child Development*, **73**, 1656-1677.
- SCHLOTTMANN, A., & ANDERSON, N. H. (1993). An information integration approach to phenomenal causality. *Memory & Cognition*, **21**, 785-801.
- SCHLOTTMANN, A., & SHANKS, D. (1992). Evidence for a distance between judged and perceived causality. *Quarterly Journal of Experimental Psychology*, **44A**, 321-342.
- SCHOLL, B. J. (2001). Objects and attention: The state of the art. *Cognition*, **80**, 1-46.
- SCHOLL, B. J., & NAKAYAMA, K. (2002). Causal capture: Contextual effects on the perception of collision events. *Psychological Science*, **13**, 493-498.
- SCHOLL, B. J., & NAKAYAMA, K. (2004). Illusory causal crescents: Misperceived spatial relations due to perceived causality. *Perception*, **33**, 455-469.
- SCHOLL, B. J., & TREMOULET, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, **4**, 299-309.
- WATANABE, K., & SHIMOJO, S. (2001). When sound affects vision: Effects of auditory grouping on visual motion perception. *Psychological Science*, **12**, 109-116.
- WEIR, S. (1978). The perception of motion: Michotte revisited. *Perception*, **7**, 247-260.
- WHITE, P. A. (1995). *The understanding of causation and the production of action*. Hillsdale, NJ: Erlbaum.
- WHITE, P. A. (in press). Visual causal impressions in the perception of several moving objects. *Visual Cognition*.
- WHITE, P. A., & MILNE, A. (1994). *Camouflaging the launching effect*. Unpublished manuscript.
- WHITE, P. A., & MILNE, A. (1997). Phenomenal causality: Impressions of pulling in the visual perception of objects in motion. *American Journal of Psychology*, **110**, 573-602.
- WHITE, P. A., & MILNE, A. (1999). Impressions of enforced disintegration and bursting in the perception of collision events. *Journal of Experimental Psychology: General*, **128**, 499-516.
- YELA, M. (1952). Phenomenal causation at a distance. *Quarterly Journal of Experimental Psychology*, **4**, 139-154.

NOTES

1. Because this basic phenomenon, like many of the effects discussed in this article, is inherently dynamic, we encourage readers to view dynamic animations of most of the conditions and figures included in this article, so that they can experience the resulting percepts. These anima-

tions can be viewed from the following Web site: <http://www.yale.edu/perception/causal-grouping/>.

2. This view was further supported by showing that the causal capture effect is largely attenuated by inserting only a small temporal asynchrony between the two events, so that the unambiguous launch happens a moment before the moment of complete overlap in the test event. See Scholl and Nakayama (2002, 2004) for details.

3. Scholl and Nakayama (2002) did report one effect wherein contextual information attenuated the perception of causality in the full-overlap test event. In their experiments, the test event was perceived in isolation as causal launching on only a small minority of trials. In their *single-context* condition, a single additional disk was added to the display, which simply traversed the entire screen in a single motion. This context event, in contrast to the primary causal capture condition, actually decreased the perception of causality: from 10.7% causal launching in isolation to only 5% with the single additional disk. This condition thus shows that contextual information can weaken the perception of causality *in a test event that is already very weak*, but it does not address the question of whether a similar type of attenuation could occur in a test event that is perceived as robustly causal in isolation. (Indeed, until the first author's discovery of reverse causal capture, the second author confidently maintained that such an effect would be impossible!)

4. Note that the effect of attention here, although both large and statistically significant, was not completely overwhelming. Even when the participants attended directly to the test event, they still tended to perceive less causal launching (74.8%) than when the test event was presented in isolation in Experiment 3 [87%; $t(9) = 2.12$; $p = .063$]. Thus,

the endogenous allocation of attention cannot completely wipe out the effect of the associated context event (perhaps because of some automatic processing of context, or perhaps because the context event still receives some degree of exogenous attention).

5. We do, however, continue to believe that such eccentricity effects exist. In particular, the results of this experiment are compatible with Michotte's (1946/1963, Experiment 7) fixation experiment, since our experiments differ in several important respects: Our two experiments used different degrees of eccentricity, different amounts of overlap in the test event, and different ambient viewing conditions, and our experiments were run in a dark room, using bright disks on a black background, whereas Michotte's experiments were run in a room with ambient light, using nonluminous objects (made with rotating disks viewed behind screens). The resulting difference in contrast may be important here, since our disks will be more salient when viewed in the periphery, and thus, causal perception may be more resistant to slight peripheral viewing.

6. We thank Peter White for drawing our attention to the importance of these camouflage experiments, which he and Alan Milne have also studied in unpublished work (White & Milne, 1994).

7. In other experiments not reported here, we are also exploring other ways in which attention and causality interact—for example, whether attention is required to perceive causality and whether collisions may actually capture attention.

(Manuscript received April 3, 2003;
revision accepted for publication September 22, 2003.)