Perception, 2006, volume 35, pages 385 – 399

DOI:10.1068/p5462

Perceiving causality after the fact: Postdiction in the temporal dynamics of causal perception

Hoon Choi, Brian J Scholl
Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA;
e-mail h.choi@yale.edu, Brian.Scholl@yale.edu
Received 6 August 2004, in revised form 25 March 2005; published online 6 February 2006

Abstract. In simple dynamic events we can easily perceive not only motion, but also higher-level properties such as *causality*, as when we see one object collide with another. Several researchers have suggested that such causal perception is an automatic and stimulus-driven process, sensitive only to particular sorts of visual information, and a major research project has been to uncover the nature of these visual cues. Here, rather than investigating *what* information affects causal perception, we instead explore the temporal dynamics of *when* certain types of information are used. Surprisingly, we find that certain visual events can determine whether we perceive a collision in an ambiguous situation even when those events occur *after* the moment of potential 'impact' in the putative collision has already passed. This illustrates a type of postdictive perception: our conscious perception of the world is not an instantaneous moment-by-moment construction, but rather is formed by integrating information presented within short temporal windows, so that new information which is obtained can influence the immediate past in our conscious awareness. Such effects have been previously demonstrated for low-level motion phenomena, but the present results demonstrate that postdictive processes can influence higher-level event perception. These findings help to characterize not only the 'rules' of causal perception, but also the temporal dynamics of how and when those rules operate.

1 Introduction

There is perhaps no concept more central to our understanding of the physical world than that of cause and effect. Accordingly, causation has received an immense amount of study, in both philosophy and psychology (see Sperber et al 1995). Much of this work has involved determining how we first acquire the concept of causality, since as an abstract concept it is unclear, for example, how it could ever be pointed out by ostentation. Moreover, Hume (1740/1960, 1748/1977) famously argued that interpretations of the world in terms of cause and effect were difficult to justify, given that all we can really ever experience is repeated correlations, eg between the arrival of one billiard ball and the departure of another. Despite this, the ability to interpret the world in terms of causal relationships is universal (eg Morris and Peng 1994), and is supported by intricate cognitive processing in adults (eg White 1995), young children (eg Gopnik et al 2004), and even infants (eg Leslie and Keeble 1987).

One interesting possibility is that a primitive notion of causality may arise from automatic visual interpretations of simple mechanical events. Though we normally think of perception in terms of recovering the physical structure of the world, several researchers following Michotte (1946/1963) have suggested that the *causal* structure of the world can in some cases also be directly perceived. We see a billiard-ball collision, for example, not just in terms of the cessation of motion in one ball and the onset of motion in another (as depicted in figure 1a): instead, we see one ball *cause* the other's motion. Consistent with the idea that such impressions reflect automatic visual processing, the perception of causality appears to be largely automatic, often irresistible, resistant to higher-level beliefs and intentions, and driven by highly constrained and stimulus-driven visual cues (see Scholl and Tremoulet 2000).

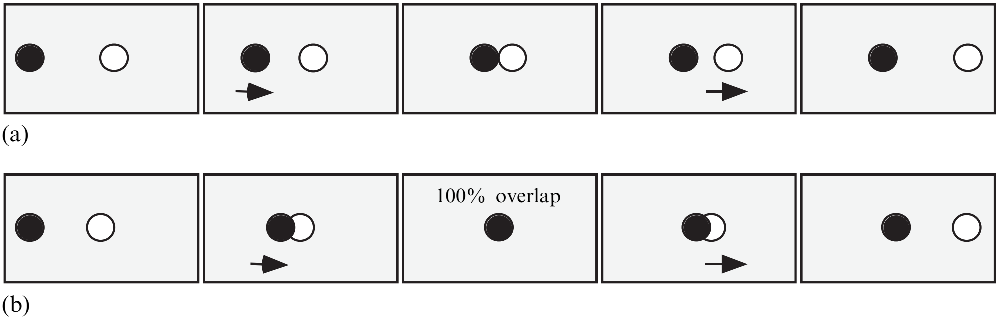

Figure 1. Two standard visual events employed in studies of causal perception. These events are presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion. (a) A 'causal launching' event, wherein one disc is seen to collide with a second disk, causing it to move. (b) A 'full-overlap' event, which observers can readily perceive as non-causal 'passing', wherein one moving object is seen to pass right over another stationary object (swapping colors).

Ever since Michotte's seminal studies (1946/1963), a great deal of research has been devoted to uncovering the nature of the visual cues which drive causal perception in simple collision events, mapping out the roles of factors such as the relative and absolute speeds of the objects, various types of spatial and temporal gaps in the objects' trajectories, the smoothness and continuity of motion, differences in the durations and angles of each object's movement, cross-modal interactions, and many other factors (eg Boyle 1960; Costall 1991; Gordon et al 1990; Guski and Troje 2003; Hubbard and Ruppel 2002; Kruschke and Fragassi 1996; Michotte 1946/1963; Michotte and Thinès 1963/1991; Natsoulas 1961; Schlottmann and Anderson 1993; Schlottmann and Shanks 1992; Scholl and Nakayama 2004; Weir 1978; White 2005a, 2005b; White and Milne 1997, 1999; Yela 1952).

1.1 *Contextual effects on causal perception*

Other recent studies have expanded the scope of this catalog of visual cues by exploring contextual effects wherein causal perception is influenced or even wholly determined by the behavior of additional objects which are not involved in the putative collision. Such studies of contextual factors have their roots in some of Michotte's demonstrations, wherein extra information added to a display can in certain cases abolish or 'camouflage' causal impressions (Michotte 1946/1963). Recent experiments have expanded on such demonstrations by showing that contextual information can also give rise to perceived causality in certain circumstances. Here, we briefly describe two such effects, which we then exploit in the present experiments to study the temporal dynamics of causal perception.

In the 'causal capture' effect (Scholl and Nakayama 2002), observers view an event which is causally ambiguous. This 'full-overlap' event (figure 1b) is similar to a canonical launching event (figure 1a) except that one disc fully occludes the other: an object, A, moves towards a stationary object, B, until they fully overlap, at which point A stops and B starts moving along the same trajectory. In isolation, this event can be readily perceived in terms of non-causal '*passing*': one object remains stationary while another passes over it. This occurs despite the salient featural differences: eg with one red object and one green object, observers can perceive a single stationary object which changes from red to green, and a single moving object which changes from green to red.

A striking effect emerges, however, when a second event—involving two additional objects, C and D—is added to the display (see figure 2). These additional objects involve a canonical collision event (figure 1a), ie an event similar to the full-overlap event, but in which D begins moving as soon as C arrives adjacent to it. When this 'contextual'

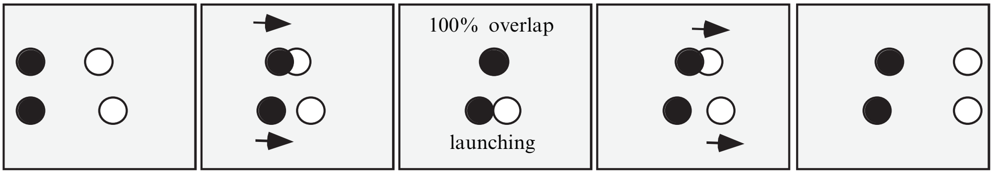

Figure 2. Depiction of the 'causal capture' phenomenon (Scholl and Nakayama 2002), presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion. Observers judge the causal status of the upper full-overlap event (which in isolation is typically seen as non-causal passing). In the presence of the temporally synchronized unambiguous 'launching' context event, the full-overlap event is now also perceived as causal launching.

event is placed in the same display as A and B (temporally synchronized such that the full overlap in A and B occurs at the same moment as the 'impact' in C and D), observers are suddenly no longer able to easily perceive A and B as non-causal passing. Instead, the unambiguous causal status of the context event (C and D) effectively 'captures' A and B, so that it too is seen as a causal launch (despite the overlap). (This then has the further effect of causing observers to perceive slightly less than 'full overlap' in the full-overlap event—as if the two objects must have been at least slightly separated in order to causally interact; see Scholl and Nakayama 2004.)

Whereas the 'causal capture' effect relies on the presence of a context event which itself is perceived in causal terms, another recent study of perceptual grouping demonstrated that the perception of causality in a full-overlap event can be induced even by the behavior of a *single* additional object which by itself yields no causal impression (Choi and Scholl 2004). In this grouping effect (figure 3a), observers again must judge the causal status of a full-overlap event (involving A and B), but a single additional object, C, is added to the display, initially vertically aligned with B. When C stays aligned with B through its motion—ie remaining stationary until A and B fully overlapped, and then moving along with B thereafter—observers again perceived the full-overlap event in terms of causal launching. This effect depends on the perceptual grouping between B and C, driven by proximity and common motion, since the full-overlap event

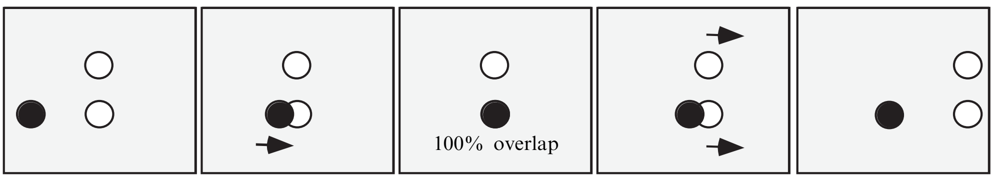

(a)

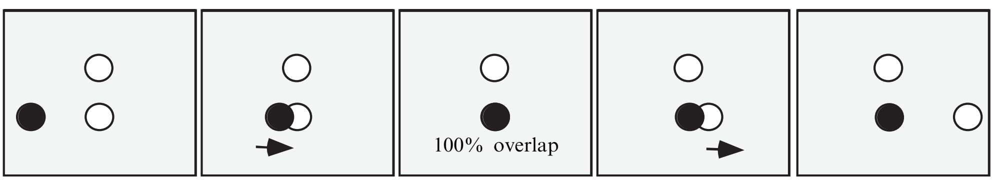

(b)

Figure 3. Effects of perceptual grouping on causal perception (Choi and Scholl 2004). Each event is presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion. (a) Observers judge the causal status of a full-overlap event. When an additional object remains grouped with the initially stationary object via proximity and common motion, observers perceive the full-overlap event as causal launching. (b) When the additional object remains stationary even after the second object begins to move, observers perceive the full-overlap event as non-causal passing.

is still readily seen as non-causal passing when C simply remains stationary throughout the event (as in figure 3b).

These contextual effects, involving causal capture and perceptual grouping, indicate that the visual cues which drive causal perception are sensitive not only to the objects involved in a potentially causal event, but also to the other objects and events located nearby. Both types of contextual events can be explained in terms of the hardwired tendency of the visual system to avoid certain types of coincidences when interpreting ambiguous visual input (eg Knill and Richards 1996; Marr 1982). In causal capture, the spatiotemporal synchronization of the two events is seen as non-accidental, 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. Similarly, in the grouping effect (figure 3a), the eventual correlated movement of the two disks is seen as non-accidental: these two discs thus form a perceptual group, which does *not* fully overlap the initially moving disc—which in turn renders the event unambiguously causal.[(1)]

1.2 *The current experiments: Postdiction and temporal dynamics*

All of the studies described above—including effects of grouping and causal capture—have focused primarily on *what* visual cues influence causal perception. Here, in contrast, we focus directly on the temporal nature of causality, by asking *when* certain visual cues can affect the perception of causal launching versus non-causal passing in ambiguous displays. It is particularly striking that no previous studies have systematically explored the temporal nature of causal perception in this way, since causality is inherently dynamic and spatio*temporal*. Yet, while many studies have manipulated extrinsic temporal factors—such as the influence of temporal gaps, eg pauses of various durations before the second motion—none has systematically varied the timing of other visual cues to causality.

Here, we turn to such temporal questions for three reasons. First, we seek to determine the breadth of the temporal window over which cues can affect the perception of causality. In our previous work we have investigated the spatial analogue of this question, eg studying how near a context object needs to be in order to influence causal perception (Choi and Scholl 2004, experiment 2c). We now ask a similar question, about how near *in time* visual cues must be to induce such effects. Our second motivation for studying the temporal dynamics of causal perception in this way is the fact that our previous studies have suggested that temporal synchrony is an especially important factor in driving some aspects of causal perception (Choi and Scholl 2004; Scholl and Nakayama 2002). This view is theoretically motivated, since the degree of 'coincidence' in contextual effects (as discussed above) is directly correlated with the degree of temporal synchrony in the motions. However, these previous studies have not systematically determined just how much synchrony is required for such effects, though we know that considerable asynchronies will obliterate effects such as causal capture (Scholl and Nakayama 2002, experiment 3).

A final reason for studying such temporal effects was discovered early on in our pilot studies for this project, when we found that cues occurring even *after* the moment of full overlap in the judged event can still determine whether that event is seen as causal or not. This is seemingly paradoxical, since, after the moment of full overlap, the visual system must interpret the ensuing motion as belonging to one of the objects,

[(1)] The notion of coincidence avoidance can even be appealed to in order to explain causal perception in simple launching events in isolation (figure 1a). Here the spatial contiguity and temporal contingency between the two objects and their movements is seen as non-accidental, and we thus perceive two objects but only a *single* motion which is transferred between them (Michotte 1946/1963). In a non-causal interpretation, in contrast, the second motion is unrelated to the sudden end of the first motion and the arrival of the first object—implying that the spatial contiguity and temporal synchrony are accidental.

and which object is bound to this motion effectively determines launching versus passing. Nevertheless, cues which occur shortly after this critical moment can apparently still cause the visual system to change this binding 'after the fact', to influence the immediate past in conscious awareness. In other words, even after you should already be perceiving non-causal passing, a visual cue can cause your percept to involve causal launching. We report two instances of this intriguing effect below, based on effects of perceptual grouping (experiment 1) and causal capture (experiment 2), and we relate this finding to other 'postdictive' effects in visual perception in section 4.

2 Experiment 1: Perceptual grouping

To begin exploring the temporal dynamics of causal perception, we first manipulated the timing of perceptual grouping in a simple contextual effect. As described above, and depicted in figure 3, observers always judged the causal status of a full-overlap event (involving objects A and B), in the presence of a single additional object, C. In our previous studies, we showed that correlated motion of B and C results in the percept of A causally launching B. This effect is especially well-suited to temporal manipulation, since adjusting the movement of C in time should directly affect the degree of grouping (ie the degree of 'commonality' in the common motion) between B and C, which should in turn influence the causal perception. (Whenever the grouping cue is insufficient to group B and C, then A still 'fully overlaps' with B, which is consistent with non-causal passing; in contrast, when B and C are grouped, A is only partially overlapping this group.)

In our previous work, B and C always moved in perfect synchrony (Choi and Scholl 2004). Here, in contrast, C again begins as a stationary disc aligned with B, but it then begins moving either before or after B begins its motion, with the degree of temporal offset systematically varied in both directions. This manipulation could yield three possible patterns of results (figure 4). First, the motion of C may give rise to causal perception only when it is perfectly synchronized with B. Thus, even the smallest temporal offset between B and C—which in our experiment consisted of the movement for a distance equal to the radius of a single object—could obliterate causal perception (figure 4a). Second, this function may be graded but asymmetric: the grouping effect could be robust in the face of small offsets when C begins moving *before* B, but this sensitivity may end as soon as B begins moving. At this point B's motion must be perceived as the motion of some object, either the initially moving object (which would

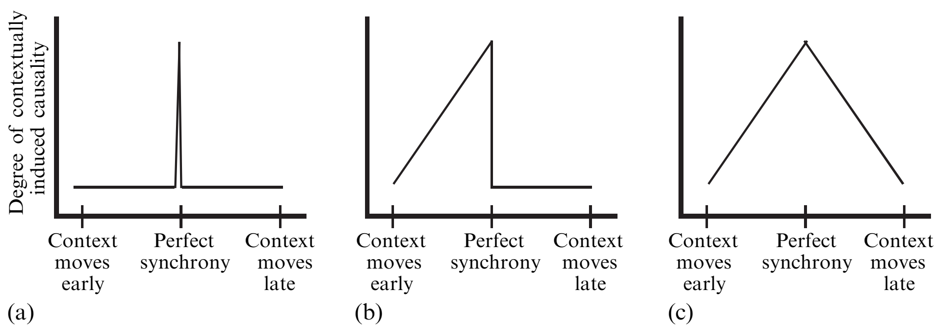

Figure 4. Three possible outcomes of experiment 1 (see text for details). (a) The common motion of the context object, C, could induce causality in a full-overlap event (involving A and B) only when the motion of B and C is perfectly synchronized. (b) The grouping effect could be robust in the face of small offsets when C begins moving *before* B (ie *pre*dictive effects), but this sensitivity could end as soon as B begins moving (after which it is too late for C's motion to affect the interaction of A and B, which is already complete). (c) If causal perception involves postdictive processes, C's motion may give rise to causal perception even when it occurs *after* the moment of full overlap in A and B.

imply non-causal passing) or the initially stationary object (which would imply causal launching). As depicted in figure 4b, this view thus predicts a sharp discontinuity, such that any motion of C occurring *after* B's motion is too late to affect the interaction between A and B (which is already complete).

To foreshadow our results, however, we can also identify a third possibility, in which the grouping effect is resistant to small temporal offsets *both* before and after the second motion begins. This view predicts a relatively symmetric function (figure 4c), and would imply the existence of a type of 'postdictive perception'. Effects of C's motion after A and B have already overlapped would seem to be too late to affect the perception of an interaction between A and B (which must by that time already be finished). This is true only if conscious awareness is a moment-by-moment construction, however. If, instead, information was integrated over short temporal windows, then perhaps C's motion could affect the perception of A and B during their overlap, even shortly 'after the fact'.

The aim of this experiment was to distinguish among these three possibilities, and thereby tell us *when* at least one visual cue can affect perceived causality.

2.1 *Method*

2.1.1 *Participants.* Fourteen naive Yale students participated for course credit or a small monetary payment. All observers had normal or corrected-to-normal acuity and normal color vision.

2.1.2 *Apparatus.* All displays were presented on a Macintosh iMac computer. Observers were positioned approximately 50 cm from the monitor, without head restraint, so that the display subtended 34.2 deg by 26 deg of visual angle. The display refreshed at 117 Hz, and motion was always perceptually smooth. All displays were constructed with custom software written with the use of the VisionShell graphics libraries (Comtois 2003). All experiments were conducted in a darkened room.

2.1.3 *Stimuli.* Each trial involved two components: one visual event consisting of two disks (A and B), and one additional context disc (C). Each disc subtended 3.18 deg, and was drawn in a solid color (see below) on a black background. Motion was always in the horizontal plane, since the perception of causality is weaker in other orientations (Michotte 1946/1963).

A and B always moved in a 'full overlap' event, as follows. The initial display consisted of two discs (one red, one green) in the same horizontal plane, with centers of the discs positioned 9.54 deg above the lower display border. Disc A started out near either the right or left edge of the display (randomly chosen on each trial), so that its most extreme edge was 2.12 deg from the display border. Disc B was always initially drawn near the center of the display, such that the nearest edges of A and B were separated by 10.07 deg. After 100 ms, A began moving at 17.65 deg s^{-1} toward B. When the two objects fully overlapped (with A always drawn on top of B; cf Scholl and Nakayama 2004), 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 2.12 deg from the display border). This entire event lasted 1.5 s.

Disc C was the same size and color as B, and was initially positioned directly above B, so that their nearest edges were vertically separated by 3.18 deg. At some point during the trial, C began moving at the same speed and in the same direction as B. C began its motion at one of nine possible moments. One of these moments was perfectly synchronous with the onset of B's motion, with four possible earlier motion onsets and four possible later motion onsets. These additional onset times were divided equitemporally in steps corresponding to the duration it took B to move one radius length (approximately 90 ms). C's motion then continued toward the other side of the display for 750 ms, at which point the display remained stationary for 100 ms before disappearing.

2.1.4 *Procedure and design.* In each trial, observers simply viewed the display until the objects disappeared, then reported via one of two key-presses whether they had perceived the visual event as a causal *launch* (in which A collided with B, causing its motion) or as a non-causal *pass* (in which A simply traversed the entire length of the display, changing colors, and passing over a stationary B, which also changed color). Observers first completed a small number of practice trials, the data from which were not recorded. During these practice trials, observers were familiarized with the particular conditions employed in that experiment, until they could reliably perceive both launching and passing in the full-overlap event. Observers then completed 20 trials of each of the nine possible conditions (corresponding to the moment when C began its motion), for a total of 180 trials; all trials were presented in a random order, unblocked. The dependent measure in this experiment is thus the percentage of perceived 'causal launching' which is reported for each possible temporal offset of C's motion.

2.2 *Results and discussion*

The percentage of trials for each possible temporal offset that were perceived as causal launching is depicted in figure 5. Clearly the temporal offset of the contextual motion had a large and systematic impact on the perception of causality. Contextual motions which began *earlier* than the moment of full overlap in the judged event (depicted in the left half of figure 5) had a straightforward influence on the results, constituting *predictive* effects wherein visual information biases the interpretation of later events. The critical aspect of these results, however, involves the conditions wherein the contextual motion begins *after* the moment of full overlap (ie the right half of figure 5). As is clear from this graph, our results are most consistent with the 'postdictive' interpretation (see figure 4c): a significant degree of causal launching was perceived even when disc C began moving after the two objects in the judged event had already overlapped. Indeed, as can be readily seen in this graph, relatively more causal launching was perceived for onsets which followed the moment of full overlap (the gray shaded region) than for those which preceded the moment of full overlap (the unshaded region).

These impressions were verified through statistical analyses. A single-factor repeated-measures ANOVA revealed a significant effect of the temporal onset of the contextual motion ($F_{8,104} = 7.69$, $p < 0.001$). Specific comparisons then verified that even contextual motions which began after the moment of full overlap could boost causal perception. For each of these tests, we compared the degree of perceived causal launching at each possible onset moment with a 'baseline' measure, operationalized here as the degree of launching perceived when the contextual object began moving at the latest possible moment (ie after the second object in the judged event had already moved 4 radii beyond full overlap, depicted as the rightmost data point in figure 5). We observed elevated degrees of causal launching due to the contextual motion when its onset occurred after the motion in the judged event had already proceeded over a distance equal to 1 radius or 2 radii—but not 3 radii—beyond the full overlap (1 radius versus 4 radii: $t_{13} = 4.29$, $p = 0.001$; 2 radii versus 4 radii: $t_{13} = 2.34$, $p = 0.036$; 3 radii versus 4 radii: $t_{13} = 1.61$, $p = 0.132$).

These results indicate a type of postdictive process in causal perception, and also illustrate its limits. Even when the two objects in the full-overlap event had moved until they were completely separated before the onset of the contextual motion, elevated degrees of perceived causal launching were still observed. In this situation, however, you cannot see the post-overlap motion without seeing that motion as bound to one of the objects: either the initially stationary disc (indicating launching) or the initially moving disc (indicating passing). The motion here, in other words, is unavoidably perceived as the motion *of a particular object*. Moreover, we know that, when there is

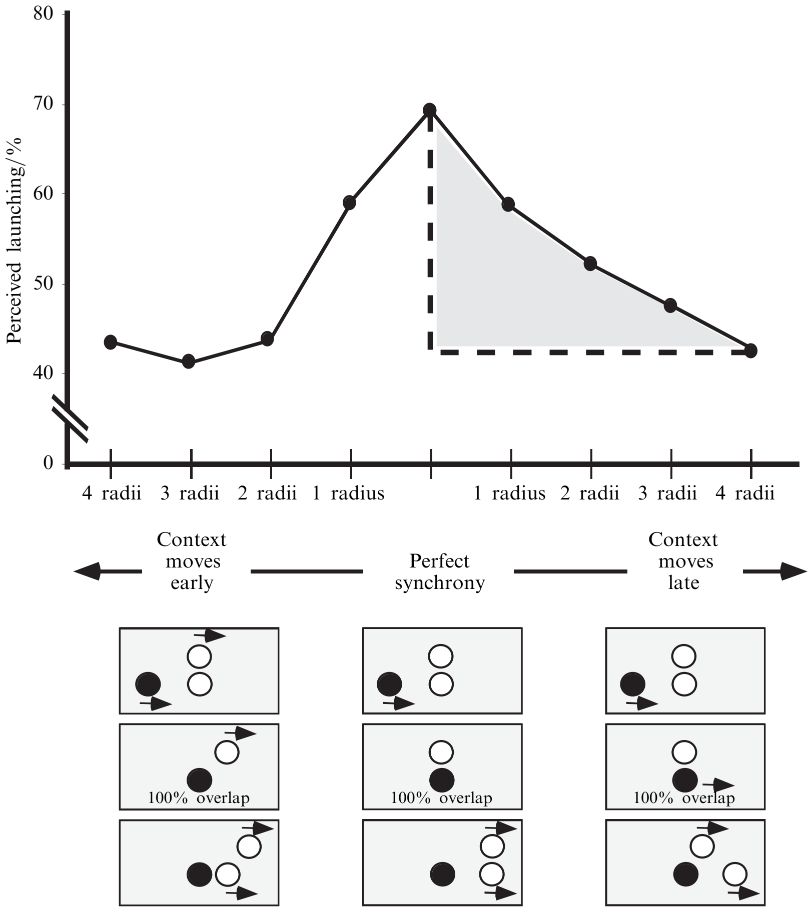

Figure 5. The percentage of perceived causal launching for each temporal offset in experiment 1. The context motion affected causal perception even when it was temporally asynchronous with the other motions. The dashed line indicates the result predicted by a view in which conscious awareness of the event is a moment-by-moment construction. The gray shaded area thus represents postdictive effects, in which the onset of the contextual motion yields a perceived collision even after the moment of full overlap has already passed.

no contextual movement (or when the contextual movement occurs extremely late), observers perceive the judged event as non-causal passing in a majority of trials, which implies that they bind the initial post-overlap movement to the initially moving disc. (This is what gives rise to the perception of a single moving object passing over a stationary object.) Thus, by the time the two objects are separated (post-overlap), this binding between the motion and an object must already be completed, with the motion bound on most trials to the initially moving disc.(2)

(2) An alternative explanation for these results could appeal to the possibility that observers simply didn't notice the temporal offset or the two distinct phases of motion at all. (We thank an anonymous reviewer for raising this possibility.) In fact, however, even the smallest temporal offsets used in this experiment were extremely salient. We verified this in a control experiment in which eight observers watched 3 trials each of the three conditions (in a randomized order): (i) perfect synchrony, (ii) the smallest temporal asynchrony wherein the context moves early, and (iii) the smallest temporal asynchrony wherein the context moves late. On each trial, observers simply had to discriminate whether the motions were synchronous or not—and were 95.83% accurate in these judgments, with only one of the eight observers performing less than perfectly.

Despite this, however, contextual motion onsets that occurred even after the two discs had completely separated (after 180 ms) were able to influence the resulting percept 'after the fact', so that the motion was bound to the initially stationary disc, yielding increased causal launching percepts. By the time the two discs had become separated by a full radius of blank space, however—ie 3 radii of motion (or 270 ms) after the full overlap—the contextual motion onset could no longer influence the resulting percept. Thus, postdictive processes appear to be able to reach 'into the past' in order to influence the construction of high-level event percepts, but the temporal windows over which such processes operate are limited to around 400 ms in this context, 200 ms before and after the moment in question.

3 Experiment 2: Causal capture

In order to test the generality and robustness of the postdictive result obtained in experiment 1, and to see whether such effects would also appear in more complicated displays, we also explored the temporal dynamics of the 'causal capture' effect described above in section 1.1. In the original causal-capture experiments, the 'impact' in the unambiguous causal launching event occurred at either the same moment as the full overlap in the judged full-overlap event (Scholl and Nakayama 2002, experiment 1; see figure 2) or up to 200 ms *before* the full overlap (experiment 3). These authors did not test for effects of contextual collisions which occurred *after* the full overlap, however, since it seemed intuitively impossible that such conditions could yield causal capture. Inspired by the results of experiment 1, however, we now tested for effects of causal capture when the moment of the contextual collision was systematically varied to occur both before and after the full overlap in the judged event.

3.1 *Method*

The procedure for this experiment was identical to that in experiment 1 except as noted below. Twelve naive Yale students participated, none of whom had participated in experiment 1. The judged event was identical to the full-overlap stimulus from experiment 1, but the context stimulus now involved an unambiguous launching event rather than a single moving object. This new context event was displayed in the same horizontal plane as the contextual object from experiment 1, and was identical to the full-overlap event, except that the first disc (C) stopped moving, and the second disc (D) started moving, as soon as those two discs became adjacent near the center of the display. The context discs were the same size as those from the full-overlap event, and the colors of the two initially moving discs (A and C) matched, as did the colors of the two initially stationary discs (B and D). Disc C began vertically aligned above disc A. The horizontal position of disc D was initially slightly to the right of the horizontal position of disc B, such that C and D were together centered in the display at the moment at which they became adjacent (as in figure 2).

The independent variable in this experiment was the temporal offset between the moment of full overlap between A and B, and the moment at which C arrived adjacent to D. These manipulations were made simply by beginning C's motion at different times, to produce either perfect temporal synchrony between these two events, or eight levels of asynchrony (four steps in which the context motion started 'early' and four in which it started 'late'), with the same timings as in experiment 1. Observers completed 20 trials of each of the nine possible timings, for a total of 180 trials; all trials were presented in a random order, unblocked.

3.2 *Results and discussion*

The percentage of trials for each possible temporal offset that were perceived as causal launching is depicted in figure 6. As is clear from this graph, our results are again most consistent with the 'postdictive' interpretation (see figure 4c): a significant degree of causal launching was perceived even when the 'impact' in the contextual collision event was asynchronous with the moment of full overlap in the judged event, and this was true both when the contextual collision preceded the moment of full overlap (*pre*diction) and when it followed the moment of full overlap (postdiction).

These impressions were verified through statistical analyses. A single-factor repeated-measures ANOVA revealed a significant effect of the temporal synchrony of the context event ($F_{8,88} = 6.10$, $p < 0.001$). Specific comparisons then verified that even contextual collisions which occurred after the moment of full overlap could induce causal capture. For each of these tests, we compared the degree of perceived causal launching at each possible contextual-collision moment with a 'baseline' measure, operationalized as the degree of launching perceived when the contextual collision occurred at the latest possible moment (ie after the second object in the judged event had already moved 4 radii beyond full overlap, depicted as the rightmost data point in figure 6). We observed significant degrees of causal capture when the contextual collision occurred after the motion in the judged event had already proceeded over a distance equal to

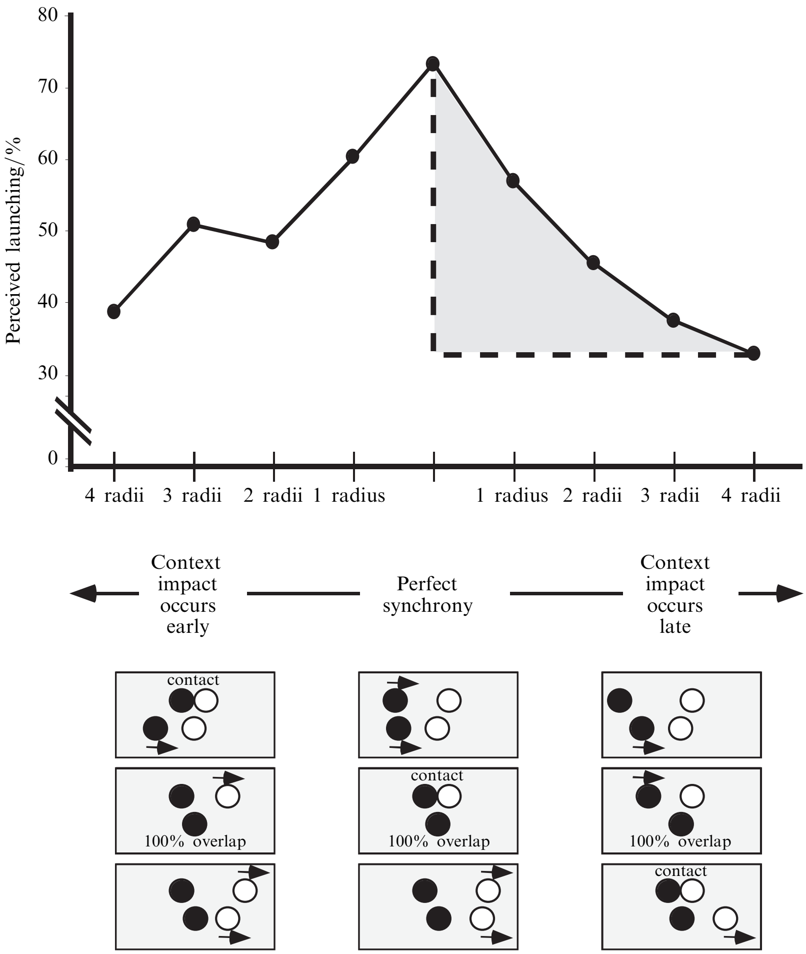

Figure 6. The percentage of perceived causal launching for each temporal offset in experiment 2. Causal capture was observed even when the contextual 'impact' was temporally asynchronous with the full overlap in the judged event. The dashed line indicates the result predicted by a view in which conscious awareness of the event is a moment-by-moment construction. The gray shaded area thus represents postdictive effects, in which the contextual impact yields a perceived collision in the full-overlap event even after the moment of full overlap has already passed.

1 radius or 2 radii beyond the full overlap, though this effect with 3 radii was only marginal (1 radius versus 4 radii: $t_{11} = 2.37$, $p = 0.037$; 2 radii versus 4 radii: $t_{11} = 3.58$, $p = 0.004$; 3 radii versus 4 radii: $t_{11} = 2.03$, $p = 0.067$).

By the same logic as employed in experiment 1, these results thus illustrate a significant degree of postdictive processing in the perception of high-level events. These results indicate that such effects occur in multiple types of contextual effects, including more complicated displays with multiple events such as the 'causal capture' display.(3)

4 General discussion

Our goal in this project was to explore the temporal dynamics of a type of high-level event perception, to determine not only *what* visual cues affect perceived causality, but also *when* those cues affect perceived causality. Our key discovery, replicated in two experiments, was a type of postdictive effect. In two experiments, observers always reported whether they perceived a 'full overlap' event as causal launching or as non-causal passing, and we explored the influence on these percepts of two representative types of contextual stimuli: the effect of grouping by proximity and common motion (experiment 1: Choi and Scholl 2004), and the causal capture effect (experiment 2; Scholl and Nakayama 2002). In both cases, we varied the temporal asynchrony between the judged full-overlap event and the contextual stimuli. (For the grouping displays, this synchrony concerned the moment of full overlap in the judged event versus the moment at which the contextual object began moving; see figure 5. For the causal capture display, this synchrony concerned the moment of full overlap in the judged event versus the moment of the contextual 'collision'; see figure 6.) We observed contextual effects (as higher percentages of perceived causal launching) in both experiments even with temporal asynchronies, and even when the critical moment in the contextual event occurred *after* the moment of full overlap in the judged event.

4.1 *Causality and postdiction*

The results of both experiments illustrate a type of postdictive perception. The motion which continues after the full overlap in the judged event must be seen as belonging to one of the two objects as soon as it begins, but our results indicate that contextual events which occur after this moment can 'influence the past' to change *which* object the motion is bound to; in particular, to re-bind the motion to the initially stationary object (yielding perceived launching), even after the motion must have been already initially bound to the initially moving object (so as to yield perceived passing in the absence of a contextual event). This seemingly paradoxical effect indicates that our conscious perception of the world is not an instantaneous moment-by-moment construction, but, rather, is formed by integrating information presented within short temporal windows, so that new information which is obtained can effectively influence the immediate past in constructing our conscious awareness.

This initial binding must have occurred by at least 90 ms after the overlap (which was the smallest asynchrony used in our study) on logical grounds, since by this point the post-overlap motion was readily apparent (as was the asynchrony itself; see footnote 2). In fact, however, there is independent evidence which suggests that this initial binding—ie the interpretation which is eventually altered by the postdictive processing—is in place by only 25 ms after the moment of the overlap. In particular, the original

(3) Note that the results in the two experiments (shown in figures 5 and 6) were extremely similar. Indeed, the reported degrees of causal launching did not differ across the two experiments for any of the temporal offsets wherein the contextual event occurred after the moment of full overlap (all ps > 0.32). Because experiment 1, unlike experiment 2, did not in fact involve a contextual *event*, all of these effects may have been driven by forms of perceptual grouping (see Choi and Scholl 2004, for discussion).

demonstration of 'causal capture' included a condition in which only a small portion of the context event was displayed, temporally 'centered' around the moment of full overlap in the judged event (Scholl and Nakayama 2002, experiment 2). Even when the context event was present for only 50 ms in total, and thus disappeared only 25 ms after the full-overlap, this was still sufficient to influence the percept of the judged event, changing it from passing (as was likely without any context event) to causal launching. Thus, while the present experiments do not provide direct evidence for exactly when the initial interpretation was constructed, it may have existed for a considerable duration before the postdictive processing driven by the subsequent motion 90 ms later.

Note that what is being postdictively influenced in these situations is not a previously generated conscious *percept*, but simply an early (and not-yet-conscious) visual representation generated before the contextual event occurred. In this situation, postdictive processing entails that conscious perception involves a type of temporal lag, such that events occurring after time *T* (by a hundred milliseconds or so) can affect the conscious perception of *T*, which is still in the process of being constructed in the visual system. We suggest that such processing may actually be quite common in visual perception, although it can only be recognized as such when ambiguous stimuli, such as the full-overlap event employed here, are used. Faced with an ambiguous event, an initial interpretation is perhaps weak enough to be overridden by more definitive information which occurs slightly later. In unambiguous events, in contrast (eg an unambiguous collision), the evidence for the initial visual hypothesis is perhaps too strong to be overridden by information which occurs later. In these situations, there may still be postdictive *processing* in the visual system—ie the visual system may still be integrating information over brief temporal windows—but the postdiction may not have any perceptible effect, either (i) because it is too weak to do so, or (ii) because it continues to *support* the initial interpretation.

4.2 *Related evidence*

Our results are consistent with previous demonstrations of postdictive processing in lower-level types of motion perception. To take a common example, postdictive processing is also implied in almost any demonstration of apparent motion (eg Kolers and von Grünau 1976; Ramachandran and Anstis 1986). When we see a disc move by apparent motion from location A to B, for example, we can in some cases see the motion traverse the intermediate positions. However, this too seems paradoxical: the visual system cannot infer motion between A and B until B has already occurred—otherwise how would the visual system know in which direction A is going to move? By the time that the flash at B has occurred, however, it should be too late to perceive motion leading up to B. Thus, the perception of apparent motion itself indicates a type of postdictive processing.

Similar examples have recently been studied for more specialized types of apparent motion. In illusory line motion, for example, a line is presented instantaneously. When preceded by a cue at one end of the line, however, observers may perceive the line as being drawn progressively, emanating from the cued end (eg Hikosaka et al 1993a, 1993b). As an illustration of postdictive processing, another recent study demonstrated that additional events which occur *after* the appearance of the line can affect the strength of the illusory line motion, even canceling it out completely (Eagleman and Sejnowski 2003). Results consistent with such postdictive processing in other domains include studies of motion-induced blindness (eg Mitroff and Scholl 2004), the flash-lag effect (eg Eagleman and Sejnowski 2000), and backward masking (eg Bachmann 1994).[4]

[4] Popular press accounts (eg Sacks 2004) suggest that dreams may also incorporate postdiction: myoclonic jerks of the body are frequently incorporated into dreams, 'antedated' so that the dream's narrative provides an explanation leading to the jerk.

Note that the logic employed in our experiments was exactly that involved in these previous studies. In the illusory line-motion study of Eagleman and Sejnowski (2003), for example, the authors employed two key stimuli. In the control stimulus, illusory line motion was perceived at the moment the line appeared. (This is analogous to our case when the context event moves far too late, or not at all, and non-causal passing is perceived.) The visual representation of (illusory) motion which ultimately supported this percept must have been generated when the line itself appeared (just as the relevant object/motion binding in our experiments must have been represented as soon as the motion began). However, when an additional context event occurred after the initial presentation of the line, this earlier representation (of motion, which must have existed) was overridden, such that observers no longer perceived illusory line motion. (Similarly, in our study, when an additional context motion occurred after the moment of full overlap, the non-causal 'passing' representation—ie the binding of the motion to the initially moving object, which must have existed—was overridden such that observers no longer saw non-causal passing.)

In fact, even some previous results involving the perception of causality can be interpreted in postdictive terms, by exactly this logic. For example, Michotte (1946/1963) and Natsoulas (1961) both report that the speed of the second motion in an unambiguous causal collision can influence the strength of the resulting causal percept. This result was not linked to the idea of postdictive processing per se, but it seems readily interpretable in such terms: one might think that the strength of a causal percept must be fixed at the time that the collision occurs, whereas the speed of the ensuing motion must necessarily be perceived slightly later. More recent research has demonstrated similar effects in ambiguous 'bouncing versus streaming' events (eg Bertenthal et al 1993; Michotte 1946/1963, experiments 24 and 97; Mitroff et al 2005), which at most speeds are perceived in causal terms, though not as causal *launching* per se. In such events, a sudden auditory stimulus that occurs at the moment of potential 'impact' can strengthen the bouncing percepts (Sekuler et al 1997; see also Watanabe and Shimojo 2001). Here, the authors focused primarily on the effect of audio-visual interaction, but their results also implicate some postdiction, since even noises presented 150 ms *after* the overlap of the two items increased perceived bouncing.

In all of these examples, as in the experiments reported here, it appears that the visual system constructs the conscious percept of an event by taking into account even visual information which arrives a short while after that event. In our experiments, the systematic manipulation of this timing suggests that such postdiction is limited to a short temporal window of less than 200 ms (which, of course, may also be combined with an equal or longer *pre*dictive window extending in time before the relevant moment). Allowing for some additional time for the relevant representations to fully decay, these temporal windows may correspond to the duration of visual iconic storage (White 1988). This temporal cutoff may represent the most adaptive balance between two constraints. On the one hand, the longer the temporal window over which the visual system integrates information to determine our conscious percepts, the more accurate these percepts might be, and the less susceptible to misleading cues on a temporally local scale. On the other hand, while our conscious perception can perhaps suffer a hundred milliseconds or so of temporal delay in many cases without undue consequences (or, indeed, without us even noticing; cf Dennett and Kinsbourne 1992), longer temporal delays might render visual perception ultimately more accurate but unsuited for real-time navigation and interaction with the world.

4.3 *Causality and coincidence avoidance*

These temporal dynamics also have implications for our understanding of the underlying nature of the causal effects themselves. Above, we suggested that grouping effects and causal capture (and even simple launching events) can be explained by appeal to the tendency of the visual system to avoid coincidences (eg Knill and Richards 1996; Marr 1982). When the two events in a causal-capture display are temporally synchronized, for example, the visual system treats this as a coincidence in need of explanation, and thereby imputes causality to the otherwise non-causal full-overlap event as a way of making this relationship less accidental. Under this interpretation, the current results indicate just how much temporal synchrony is needed to 'count' as a coincidence. Apparently up to about 200 ms of temporal asynchrony can be tolerated in this process: as long as the two events occur within this temporal window, the visual system assumes they must be related. And the fact that this temporal window extends a short while into the future (in addition to the past) in causal perception demonstrates that postdictive processing is not isolated to lower-level types of motion perception, but also interacts with higher-level event perception, and thus may be a more common aspect of visual processing and conscious awareness.

Acknowledgments. For helpful conversation and/or comments on earlier drafts, we thank Erik Cheries, Steve Mitroff, Peter White, and two anonymous reviewers. BJS was supported by NSF #BCS-0132444.

References

Bachmann T, 1994 *Psychophysiology of Visual Masking* (Commack, NY: Nova Science)
Bertenthal B, Banton T, Bradbury A, 1993 "Directional bias in the perception of translating patterns" *Perception* **22** 193 – 207
Boyle D, 1960 "A contribution to the study of phenomenal causation" *Quarterly Journal of Experimental Psychology* **12** 171 – 179
Choi H, Scholl B J, 2004 "Effects of grouping and attention on the perception of causality" *Perception & Psychophysics* **66** 926 – 942
Comtois R, 2003 *VisonShell PPC* [Software libraries] (Cambridge, MA: author)
Costall A, 1991 "Phenomenal causality", in *Michotte's Experimental Phenomenology of Perception* Eds G Thinès, A Costall, G Butterworth (Hove, Sussex: Lawrence Erlbaum Associates) pp 51 – 64
Dennett D C, Kinsbourne M, 1992 "Time and the observer: The where and when of consciousness in the brain" *Behavioral and Brain Sciences* **15** 183 – 247
Eagleman D M, Sejnowski T J, 2000 "Motion integration and postdiction in visual awareness" *Science* **287** 2036 – 2038
Eagleman D M, Sejnowski T J, 2003 "The line-motion illusion can be reversed by motion signals after the line disappears" *Perception* **32** 963 – 968
Gopnik A, Glymour C, Sobel D, Schulz L, Kushnir T, Danks D, 2004 "A theory of causal learning in children: Causal maps and Bayes-Nets" *Psychological Review* **111** 1 – 31
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, 2003 "Audiovisual phenomenal causality" *Perception & Psychophysics* **65** 789 – 800
Hikosaka O, Miyauchi S, Shimojo S, 1993a "Focal vision attention produces illusory temporal order and motion sensation" *Vision Research* **33** 1219 – 1240
Hikosaka O, Miyauchi S, Shimojo S, 1993b "Voluntary and stimulus-induced attention detected as motion sensation" *Perception* **22** 517 – 526
Hubbard T L, Ruppel S E, 2002 "A possible role of naïve impetus in Michotte's 'launching effect': Evidence from representational momentum" *Visual Cognition* **9**(1/2) 153 – 176
Hume D, 1740/1960 *A Treatise of Human Nature* (Oxford: Clarendon Press, 1960)
Hume D, 1748/1977 *An Enquiry Concerning Human Understanding* (Indianapolis, IN: Hackett, 1977)
Knill D, Richards W, 1996 *Perception as Bayesian Inference* (Cambridge: Cambridge University Press)
Kolers P, Grünau M von, 1976 "Shape and color in apparent motion" *Vision Research* **16** 329 – 335
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* (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 441 – 446
Leslie A M, Keeble S, 1987 "Do six-month-old infants perceive causality? *Cognition* **25** 265 – 288
Marr D, 1982 *Vision* (New York: W H Freeman)
Michotte A, 1946/1963 *La Perception de la Causalité* (Louvain: Institut Supérieur de Philosophie, 1946 [English translation of updated edition by T Miles and E Miles *The Perception of Causality* (New York: Basic Books, 1963)]

Michotte A, Thinès G, 1963/1991 "La causalité perceptive" *Journal de Psychologie Normale et Pathologique* **60** 9 – 36, 1963 [Translated and reprinted in *Michotte's Experimental Phenomenology of Perception* Eds G Thinès, A Costall, G Butterworth (Hillsdale, NJ: Lawrence Erlbaum Associates, 1991) pp 66 – 87]
Mitroff S R, Scholl B J, 2004 "Seeing the disappearance of unseen objects" *Perception* **33** 1267 – 1273
Mitroff S R, Scholl B J, Wynn K, 2005, "The relationship between object files and conscious perception" *Cognition* **96** 67 – 92
Morris M W, Peng K, 1994 "Culture and cause: American and Chinese attributions for social and physical events" *Journal of Personality and Social Psychology* **67** 949 – 971
Natsoulas T, 1961 "Principles of momentum and kinetic energy in the perception of causality" *American Journal of Psychology* **74** 394 – 402
Ramachandran V S, Anstis S M, 1986 "The perception of apparent motion" *Scientific American* **254**(6) 102 – 109
Sacks O, 2004 "Speed: Aberrations of time and movement" *New Yorker* 8/23/04, 60 – 69
Schlottmann A, Anderson N H, 1993 "An information integration approach to phenomenal causality" *Memory & Cognition* **21** 785 – 201
Schlottmann A, Shanks D, 1992 "Evidence for a distance between judged and perceived causality" *Quarterly Journal of Experimental Psychology A* **44** 321 – 342
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
Sekuler R, Sekuler A, Lau R, 1997 "Sound alters visual motion perception" *Nature* **385** 308
Sperber D, Premack D, Premack A, 1995 *Causal Cognition: A Multidisciplinary Debate* (Oxford: Oxford University Press)
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, 1988 "Causal processing: Origins and development" *Psychological Bulletin* **104** 36 – 52
White P A, 1995 *The Understanding of Causation and the Production of Action* (Hillsdale, NJ: Lawrence Erlbaum Associates)
White P A, 2005a "Visual causal impressions in the perception of several moving objects" *Visual Cognition* **12** 395 – 404
White P A, 2005b "Visual impressions of interactions between objects when the causal object does not move" *Perception* **34** 491 – 500
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 visual 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

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 35 2006

www.perceptionweb.com