Visual adaptation of the perception of causality

Martin Rolfs^{1,2,3,4,*}, Michael Dambacher^{5,6}, Patrick Cavanagh⁷

- ¹ Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA
- ² Laboratoire de Psychologie Cognitive, Université Aix-Marseille, 3 Place Victor Hugo, 13331 Marseille, France
- ³ Bernstein Center for Computational Neuroscience Berlin, Philippstr. 13, 10115 Berlin, Germany
- ⁴ Department of Psychology, Humboldt University Berlin, Unter den Linden 6, 10099 Berlin, Germany
- ⁵ Department of Psychology, University of Potsdam, Karl-Liebknecht-Str. 24-25, 14476 Potsdam, Germany
- ⁶ Department of Psychology & Zukunftskolleg, University of Konstanz, Box D29, 78457 Konstanz, Germany
- ⁷ Université Paris Descartes, Laboratoire Psychologie de la Perception, 45 rue des Saints-Pères, 75006 Paris, France
- * Correspondence: <u>martin.rolfs@hu-berlin.de</u>

Manuscript in press in

Current Biology

Summary

We easily recover the causal properties of visual events, enabling us to understand and predict changes in the physical world. We see a tennis racket hitting a ball and sense that it caused the ball to fly over the net; we may also have an eerie but equally compelling experience of causality if the streetlights turn on just as we slam our car's door. Both perceptual [1] and cognitive [2] processes have been proposed to explain these spontaneous inferences, but without decisive evidence one way or the other, the question remains wide open [3-8]. Here, we address this long-standing debate using visual adaptation—a powerful tool to uncover neural populations that specialize in the analysis of specific visual features [9-12]. After prolonged viewing of causal collision events called 'launches' [1], subsequently viewed events were judged more often as non-causal. These negative aftereffects of exposure to collisions are spatially localized in retinotopic coordinates, the reference frame shared by the retina and visual cortex. They are not explained by adaptation to other stimulus features and reveal visual routines in retinotopic cortex that detect and adapt to cause and effect in simple collision stimuli.

Results

We often have a strong sense of causality as events unfold, where one event apparently triggers the next. The perception of causality involves two components, one that is stimulus-based and one that is inference-based. First, to see causal structure between two events, they need to follow each other with little delay and in many cases, including collisions, they also require contact. This spatiotemporal coincidence is the stimulus-based component of perceptual causality. The second component is an inference, merging two events into one. Rather than seeing one object stopping and a second one starting on its own, there is a continuity of action that is transferred from the first object to the second. Using a visual adaptation paradigm, we tested if this inference occurs on a perceptual level, which would allow experimental access to the visual detection of causality and the parsing of events at a perceptual stage.

To measure observers' perception of causality, we displayed short animations of two *test events*, appearing either above/below or left/right of a central fixation spot (**Figure 1A**). Each test event consisted of two gray discs that at times were stationary and at other times moved across the black background. Identical but mirror-symmetrical events appeared simultaneously on both sides of the screen to discourage reflexive eye movements. In each event, one disc was initially stationary as the other one approached on a direct path from a random direction. After 80 ms, when the two discs overlapped by some amount, the moving disc froze and the stationary one took off in the same direction at the same speed (**Figure 1B**; **Movie S1**). Depending on the overlap between the two discs when the second disc starts moving (**Figure 1C**), these



Figure 1 Procedure and stimulus design.

(A) On every trial, two mirror-symmetrical test events appeared either left/right (shown here) or above/below a central fixation spot.

(B) Test event. After 80 ms of motion, Disc 1 (black) stops with some degree of overlap (here, 0.25 diameters) with Disc 2 (blue). Disc 2 then moves off in the same direction and at the same speed.

(C) Test events had one of 9 magnitudes of overlap in their central frame.

test events either appear causal or non-causal. In particular, when there is little overlap, one disc appears to causally launch the other's motion; whereas when there is large overlap, one disc appears to non-causally pass over (or under) the other which in turn remains stationary [1,13]. On every trial, observers pressed one of two buttons to indicate what they perceived, a launch or a pass. By fitting a psychometric function to each observer's data, we obtained the point of subjective equality (PSE), which captures the amount of disc overlap at which the observer is equally likely to report the event as a launch or a pass.

Adaptation to collision events, appearing causal

In our first experiment, each observer performed two blocks of trials, and we measured two PSEs each time, one for events presented left/right of fixation and one for events above/below fixation. After the first block, in which we obtained baseline PSEs for both sets of locations, observers saw an *adapting stream* of 320 'collision' (or, launch [1]) stimuli (**Figure 2A; Movie S2**)—two discs bouncing back and forth, clearly appearing causal—while maintaining fixation at the screen center. For half the observers, we showed these adapting streams only at the horizontal locations, for the other half only at the vertical locations. Adapting streams were organized in pairs with the first stimulus in each pair having a random direction of motion and the second having the opposite direction. After this adaptation phase, we obtained a second set of



Figure 2 Adaptation to causal collisions.

(A) Collision adaptation stimulus. After 80 ms of motion, Disc 1 (black) stops next to Disc 2 (blue), which then moves off.

(B) Average proportion of pass reports with cumulative Gaussian psychometric functions modeling the relationship between disc overlap and perceptual reports.

(C) Individual (black dots) and average (blue marker) changes in PSE for the adapted and unadapted test locations. Error bars are 95% confidence intervals (thick: comparison against line of equality); *p*-values are bootstrapped tests against zero. In the blue-shaded area, the decrease in Δ PSEs is larger at the adapted than at the unadapted location.

PSEs, one at the adapted locations and one at the locations where no adapting streams had appeared. Adaptation (provided it occurred) was topped-up by a further stream of 16 adapting stimuli preceding every test trial. The critical dependent variable was the change in PSE at either location. If prolonged exposure to phenomenologically causal stimuli shifts the PSE, we have evidence for adaptation.

Before adaptation, observers' PSEs were 0.63 ± 0.07 (mean \pm s.e.m.; where 1.0 indicates complete overlap) and 0.59 ± 0.08 for the (to-be) adapted and unadapted location, respectively. Following adaptation, we observed a strong shift of the





(A) Slip adaptation stimulus. Disc 1 (black) moves towards Disc 2 (blue), slips over it, and comes to stop on the other side after 130 ms of motion. Disc 2 then moves off.

- (B) Average proportion of pass reports with psychometric functions.
- (C) Individual changes in PSE. Conventions as in Figure 2C.

psychometric function to the left for test events presented at the adapted location (PSE = 0.34±0.06; **Figure 2B**, *blue* vs. *black curve*). Events that were perceptually ambiguous before adaptation (e.g., overlap of 0.625) were now judged to be non-causal passes in the vast majority of trials; events that were regularly perceived as causal before adaptation (e.g., overlap of 0.375) had now become ambiguous. Perception of test events at the unadapted location was less affected (PSE = 0.51 ± 0.06 ; **Figure 2B**, *gray* vs. *black curve*). We captured these effects by plotting the individual changes in PSE at the adapted location against those at the unadapted location (**Figure 2C**). Every observer's data point fell below the x-axis, showing that the PSE decreased at the adapted location ($\Delta PSE_{adapted} = -0.27\pm0.05$, $p < 10^{-9}$, BF > 10^{3}), resulting in a substantial adaptation effect (*blue marker*). We also observed a small adaptation effect at the unadapted location ($\Delta PSE_{unadapted} = -0.08\pm0.03$, p < 0.01, BF =

0.42). Importantly, all data points fell into the blue-shaded area, showing that the decrease in PSE was larger at the adapted than at the unadapted location ($\Delta PSE_{adapted} - \Delta PSE_{unadapted} = -0.19\pm0.03$, p < 10⁻¹⁰, BF > 10⁴). These results, therefore, uncover pronounced and spatially specific negative aftereffects of prolonged exposure to perceptually causal collision stimuli.

Adaptation to slip events, appearing non-causal

In a second step, we examined whether adaptation to other visual features of the adapting stimuli might explain the change in observers' judgments of causality. To do so, we repeated the first experiment using a 'slip' adaptation stimulus (**Figure 3A**; **Movie S3**), designed to match the collision adaptation stimulus (**Figure 2A**) in as many physical properties as possible—the number and appearance of the discs, spatiotemporal contiguity, timing and number of motion onsets, motion direction and speed, as well as the area covered by the event—without leading to perceived causality. In a slip adaptation stimulus, rather than stopping when it touches the stationary one the moving disc moves completely across and comes to stop immediately on the other side. With no delay, the initially stationary disc then starts off in the same direction, leaving the impression of two independent motions. Note that slip and collision adaptation stimuli both feature spatial as well as temporal coincidence; specifically, one object stops next to the other just as the other takes off. However, one appears causal and the other does not.

Contrary to the effects of adaptation to collisions, adaptation to slip stimuli had little or no effect on observers' perceptual reports. Before adaptation, we again observed the increase in the proportion of pass reports with increasing disc overlap of the test stimuli (Figure 3B, black), with PSEs of 0.45±0.03 and 0.41±0.05 for the (tobe) adapted and unadapted location, respectively. On average, psychometric functions shifted slightly to the right both for the adapted (Figure 3B, blue; PSE: 0.53±0.05; $\Delta PSE_{adapted} = 0.08 \pm 0.04$, p = 0.052, BF = 0.19) and the unadapted location (Figure 3B, *gray*; PSE: 0.47 ± 0.06 ; $\Delta PSE_{unadapted} = 0.07 \pm 0.04$, p = 0.076, BF = 0.15), with no difference between them ($\Delta PSE_{adapted} - \Delta PSE_{unadapted} = 0.01 \pm 0.02$, p = 0.52, BF < 10⁻¹). Individual changes in PSEs are shown in **Figure 3C**. A direct statistical comparison of these results to those from Experiment 1 again established the spatial specificity of aftereffects following adaptation to collision vs. adaptation to slip events $(-0.20\pm0.03, p < 10^{-8}, BF > 10^{3})$. We conclude that the perceptual changes after exposure to streams of collisions were not caused by adaptation to any of the visual attributes shared by collision and slip stimuli. Instead, adaptation affected visual processing of causal structure, present only in our collision stimuli.

An alternative explanation of the effect of adaptation is that it did not affect perceived causality directly but instead changed the perceived timing of the individual events which in turn reduced perceived causality [1,14-17]. Our collision and slip

stimuli were constructed to equate the timing of the two discs' motions, but to be sure, we ran a second control experiment to test potential effects of timing directly (see Supplemental Data and **Figure S1**). These supplemental results showed that we had adapted detectors of causality for this type of stimuli, not detectors of timing.

Reference frame of adaptation to collision events

In a final experiment, we determined the reference frame of adaptation to causal collision stimuli. If this adaptation of perceptual causality occurs at early stages of visual processing, we would expect adaptation to occur in a retinotopic frame of reference. In that case, the strongest aftereffects would be observed at locations falling on the same patch of the retina (and retinotopically organized brain areas). The aftereffects observed in the first experiment, however, are equally compatible with changes of perception tagged to a location in the world, possibly indicating an association of the experience with the adaptation stimulus with a location in external space. In that case, the strongest aftereffects would be observed at a fixed location in space, irrespective of the retinal location of the test. Identifying the reference frame of adaptation therefore requires changes in gaze position between adaptation and test [18].

We presented stimuli on one side of fixation and monitored eye position, ensuring that stimuli occurred at the intended retinal location (**Figure S2**). Observers made two fixation steps before a test event appeared—first away from initial fixation, then either back to fixation or to the opposite location (**Figure 4A**). Two initial blocks of test trials preceded the adaptation phase (another two blocks) where an adaptation stream of 320 collision stimuli appeared before the first and another 16 before every subsequent test trial. Depending on, (1) the correspondence of the fixation location during adaptation and test stimulus presentation and, (2) the location of the test event relative to fixation, the test either fell onto the same retinal, but different spatial location (retinotopic aftereffect, *green* in **Figure 4A**), the same spatial, but different retinal location (spatiotopic, *beige*), the same spatial and retinal location (full, *blue*), or an eccentricity-matched unadapted location (nonspecific, *gray*).

Data from four participants showed an aftereffect only if test events coincided with the adaptation stream in retinal coordinates, i.e., in the retinotopic and full aftereffect conditions (**Figure 4B**; PSEs at nonspecific location: 0.64 ± 0.02 before vs. 0.64 ± 0.02 after adaptation, BF < 10^{-1} ; full: 0.63 ± 0.03 vs. 0.48 ± 0.03 , BF > 10^{12} ; spatiotopic: 0.67 ± 0.05 vs. 0.65 ± 0.04 , BF < 10^{-1} ; retinotopic 0.64 ± 0.04 vs. 0.46 ± 0.03 , BF > 10^{2}). Accordingly, an ANOVA with test phase (before vs. after adaptation), retinal test location (adapted vs. unadapted), and fixation position (matched vs. unmatched between adaptation and test) as within-subject factors, yielded significant effects of retinal test location (*F*(1,3)=219.21, *p*<0.001) and test phase (*F*(1,3)=64.33, *p*=0.004),



Figure 4 Reference frame of adaptation to collisions.

(A) Observers fixated a spot 5° to the left or to the right of the screen center. Following the adaptation stream at screen center, observers made two fixation steps (red arrows). Subsequently, a test event occurred at 5° eccentricity, either to the left or to the right of fixation, yielding a total of four possible test locations: nonspecific (gray), full aftereffect (blue), spatiotopic (beige), and retinotopic (green).

(B) Average proportion of pass reports with psychometric functions.

as well as an interaction of these two factors (F(1,3)=67.53, p=0.004; other Fs<1.8, ps>0.27).

Discussion

The retinal specificity of the observed aftereffects argues strongly that it is a consequence of a shift in a perceptual, not a cognitive, boundary between causal and noncausal events. Previous experiments have shown that repeated exposure to or training in the categorization of causal stimuli alters the frequency of causal reports, but these shifts have been interpreted as cognitive anchoring effects [19,20]. Cognitive boundary shifts are common and may even be contingent on location in the world—

what looks like steam over a pot will look like smoke over a chimney. Never, however, will cognitive boundary shifts be specific to a particular location on our retina, independent of location in the world. Nevertheless, this retinal specificity, on its own, is not sufficient evidence that causal processing occurs at the perceptual level. The inference of causality is the end result of a chain of analyses and the adaptation of visual signals at early levels, prior to the determination of causality, could certainly produce a location-specific effect. This would be true even if the final decision stage was cognitive and non-retinotopic. Our conclusions therefore rest on our combined results of retinotopically specific adaptation and the absence of adaptation to slip events that were matched to the collision events in low-level visual signals.

Visual adaptation demonstrates the perceptual consequences of a reduction in the responsiveness of neural populations that encode primary visual features [10,12]. Using this general paradigm, we provided support for the existence of adaptable, visual neurons (or neural populations) that underlie the perception of at least one causal interaction in dynamic scenes. Stimuli that do not appear causal (including our 'slip' adaptation stimuli), leave the responses of these neurons unaffected. These neuronal populations must be located in brain areas that encode visual information in an eyecentered reference frame, as the resulting aftereffects are specific to the adapted location on the retina. Candidates for such areas are the mediotemporal area V5 and the superior temporal sulcus, both of which have eye-centered representations [21] and are part of a network involved in the perception of causal launches [22-25]. These areas also respond to other forms of meaningful motion patterns, such as biological motion [26,27]. Using adaptation, we can now examine the visual computations underlying the perception of causal structure in the visual world, including not only the routines recognizing familiar motion patterns [28], but also complex interactions involving cause and effect and potentially, animacy, and intentionality.

We have focused on one specific causal stimulus (collisions, or launches) and shown that it is adaptable. It is not yet known if these aftereffects generalize to other types of causal stimuli, but our finding takes a more important step—it isolates a visual process that merges two events into a single percept, thus parsing the continuity of action in the visual scene. This finding allows us to move phenomena that have been regarded as higher-level processes into the realm of perception, opening them to systematic study using the tools of perceptual science. In a similar sense, the discovery of amodal completion allowed us to study the integration of object structure behind an occluding surface, and the discovery of apparent motion allowed us to study the merging of two object identities at different locations and times into the motion of a single one. Both of these percepts require sophisticated inference and it is now widely agreed that perception is the locus of these advanced decisional processes. This has led to pathbreaking studies of their neural correlates [29-34]. The present findings take an equally important step for how the brain parses events and assigns causal links, which paves the way for tracking down the neural mechanisms underlying these visual processes.

Acknowledgments

M.R. was supported by the EU's 7th FP (Marie Curie IOF 235625) and the Deutsche Forschungsgemeinschaft (RO 3579/2-1); M.D. by the Deutsche Forschungsgemeinschaft (FOR868/1); P.C. by a Chaire d'Excellence Grant. We thank Marisa Carrasco and Reinhold Kliegl for making available their experimental setups.
M.R. conceived the experiments; M.R. and M.D. collected and analyzed the data; M.R., M.D., and P.C. discussed the results and wrote the paper. The authors declare no competing financial interests.

References

- 1. Michotte, A. (1963). The Perception of Causality (Oxford: Basic Books).
- 2. Hume, D. (1967). A Treatise of Human Nature (Oxford: Oxford University Press).
- 3. Rips, L.J. (2011). Causation From Perception. Persp. Psychol. Sci. *6*, 77-97.
- 4. Wolff, P. (2007). Representing causation. J. Exp. Psychol. Gen. *136*, 82-111.
- 5. White, P.A. (2009). Perception of forces exerted by objects in collision events. Psychol. Rev. *116*, 580-601.
- 6. Schlottmann, A., Ray, E.D., Mitchell, A., and Demetriou, N. (2006). Perceived physical and social causality in animated motions: spontaneous reports and ratings. Acta Psychologica *123*, 112-143.
- 7. Scholl, B., and Tremoulet, P. (2000). Perceptual causality and animacy. Trends Cogn. Sci. *4*, 299-309.
- 8. Wagemans, J., Van Lier, R., and Scholl, B.J. (2006). Introduction to Michotte's heritage in perception and cognition research. Acta Psychologica *123*, 1-19.
- 9. Clifford, C.W.G., and Rhodes, G. (2005). Fitting the Mind to the World: Adaptation and Aftereffects in High Level Vision (Oxford: Oxford University Press).
- 10. Clifford, C.W.G. et al. (2007). Visual adaptation: neural, psychological and computational aspects. Vision Res. *47*, 3125-3131.
- 11. Kohn, A. (2007). Visual adaptation: physiology, mechanisms, and functional benefits. J. Neurophysiol. *97*, 3155-3164.
- 12. Schwartz, O., Hsu, A., and Dayan, P. (2007). Space and time in visual context. Nat. Rev. Neurosci. *8*, 522-535.
- 13. Scholl, B.J., and Nakayama, K. (2002). Causal capture: contextual effects on the perception of collision events. Psychol. Sci. *13*, 493-498.
- 14. Guski, R. and Troje, N.F. (2003). Audiovisual phenomenal causality. Percept. Psychophys. *65*, 789-800.
- 15. Schlottmann, A. and Anderson, N.H. (1993). An information integration approach to phenomenal causality. Mem. Cogn. *21*, 785-801.
- 16. Schlottmann, A. and Shanks, D.R. (1992). Evidence for a distinction between judged and perceived causality. Q. J. Exp Psychol. A *44*, 321-342.

- 17. Fujisaki, W., Shimojo, S., Kashino, M., and Nishida, S. (2004). Recalibration of audiovisual simultaneity. Nat. Neurosci. *7*, 773-778.
- 18. Knapen, T., Rolfs, M., and Cavanagh, P. (2009). The reference frame of the motion aftereffect is retinotopic. J. Vis. *9*, 16.
- 19. Gruber, H.E., Fink, C.D., and Damm, V. (1957). Effects of experience on perception of causality. J. Exp. Psychol. *53*, 89-93.
- 20. Powesland, P.F. (1959). The effect of practice upon the perception of causality. Can. J. Psychol. *13*, 155-168.
- 21. Saygin, A.P., and Sereno, M.I. (2008). Retinotopy and attention in human occipital, temporal, parietal, and frontal cortex. Cereb. Cortex *18*, 2158-2168.
- 22. Blakemore, S.J. et al. (2001). How the brain perceives causality: an event-related fMRI study. NeuroReport *12*, 3741-3746.
- 23. Fonlupt, P. (2003). Perception and judgement of physical causality involve different brain structures. Brain Res. *17*, 248-254.
- 24. Fugelsang, J.A., Roser, M.E., Corballis, P.M., Gazzaniga, M.S., and Dunbar, K.N. (2005). Brain mechanisms underlying perceptual causality. Brain Res. *24*, 41-47.
- 25. Straube, B., and Chatterjee, A. (2010). Space and Time in Perceptual Causality. Frontiers Hum. Neurosci. *4*, 10.
- 26. Grossman, E. et al. (2000). Brain areas involved in perception of biological motion. J. Cogn. Neurosci. *12*, 711-720.
- 27. Saygin, A.P., Wilson, S.M., Hagler, D.J., Bates, E., and Sereno, M.I. (2004). Point-light biological motion perception activates human premotor cortex. J. Neurosci. *24*, 6181-6188.
- 28. Cavanagh, P., Labianca, A.T., and Thornton, I.M. (2001). Attention-based visual routines: sprites. Cognition *80*, 47-60.
- 29. Zhou, H., Friedman, H. S., and von der Heydt, R. (2000). Coding of border ownership in monkey visual cortex. J. Neurosci. *20*, 6594-6611.
- Weigelt, S., Singer, W., and Muckli, L. (2007). Separate cortical stages in amodal completion revealed by functional magnetic resonance adaptation. BMC Neurosci. *31*, 8:70.
- 31. von der Heydt, R., Macuda, T., and Qiu, F.T. (2005). Border-ownershipdependent tilt aftereffect. J. Opt. Soc. Am. A *22*, 2222-2229.
- 32. Muckli, L., Kohler, A., Kriegeskorte, N., and Singer, W. (2005) Primary Visual Cortex Activity along the Apparent-Motion Trace Reflects Illusory Perception. PLoS Biol. *3*, e265.
- 33. Churchland, A.K., Huang X., & Lisberger, S.G. (2007). Responses of neurons in the medial superior temporal visual area to apparent motion stimuli in macaque monkeys. J. Neurophysiol. *97*, 272-282.
- 34. von Grünau, M.W. (1986). A motion aftereffect for long-range stroboscopic apparent motion. Percept. Psychophys. *40*, 31-38.

Inventory of Supplemental Information

1. Supplemental Data

(a) Experiment 4 and Figure S1

Additional experiment showing that prolonged exposure to causal collisions adapts perceived causality, rather than perceived timing *(related to Experiment 1 and Figure 2)*

(b) **Figure S2** Distributions of eye positions during adaptation streams in Experiment 3 *(related to Experiment 3 and Figure 4)*

2. Supplemental Experimental Procedures

(c) Experimental Procedures

(d) Movie S1

Examples of test events, slowed down by a factor of ~3.3 *(related to Figure 1)*

(e) Movie S2

Examples of collision adaptation stimuli, slowed down by a factor of ~3.3 *(related to Figure 2)*

(f) Movie S3

Examples of slip adaptation stimuli, slowed down by a factor of ~3.3 *(related to Figure 3)*

3. Supplemental References

1. Supplemental Data

Exposure to causal collisions adapts perceived causality, not perceived timing

Our main experiments showed that following exposure to collision events, ambiguous events appeared to be less causal, a result that we interpreted as evidence for the adaptation of perceived causality.

An alternative explanation is an adaptation of perceived timing rather than perceived causality. In the launch stimulus that we used, the relative timing between the end of the first disc's motion and the start of the second disc's motion is critical for the perception of causality [1]. An asynchrony between the stop and the start (**Figure S1A**) reduces the perceived causality (although a slight positive delay in the onset of the second disc's motion is optimal [1,14-16]). The black line in **Figure S1B** shows this relationship.

Clearly, if prolonged exposure to collision events changes the perceived timing of starting and stopping, this would result in a decrease of causal percepts that was a secondary consequence of the adaptation of timing, rather than a direct adaptation of perceived causality. We briefly consider how our adaptation could have changed perceived timing in Experiments 1 and 3, and then present our test of this potential timing effect.

The difference between the black open marker and the blue filled marker in **Figure S1B** shows the observed decrease in the proportion of causal reports after adaptation to collision events. The adaptation stimulus always had a delay of 0 ms between start and stop of the two discs, and this moment was always simultaneous with the first contact (no overlap). Adaptation to these values might change the perceived timing in test events when they have different temporal characteristics. While our test events also had no temporal delay between stopping of the first and starting of the second disc, the timing of the first contact between the discs varied with spatial overlap (e.g., with an overlap of 0.5 diameters, the discs stopped and started simultaneously at around 25 ms after their initial contact). This difference of adaptation events potentially altered perceived timing.

To examine this possibility, we repeated Experiment 1 but manipulated the time delay of motion offset and onset in the test events rather than their spatial overlap. Any change in perceived timing following adaptation would vary the response function along the time axis by shifting or changing its width, making a different range of timings optimal for perceived causality [17]. The dotted blue line in **Figure S1B** shows one example of such a shift, with a narrowing of the width of the function (adaptation expands, repels, perceived values near the adapting stimulus). Critically, though, adaptation to timing would not predict changes in the amplitude of the response function.

Conversely, a direct adaptation of perceived causality (rather than timing detectors which in turn influence perceived causality) would predict a general decrease in perceived causality, across different delays (dashed blue line in **Figure S1B**) but no shift or change in the response function other than its amplitude.

In our control experiment, adaptation to collision events clearly decreased the amplitude of the response function but did not shift it or narrow it. Three observers (28–38 y; mean: 33.3 y; all male) completed three sessions each. In each session, they saw 12 test events of each of a set of 8 delays (between –70 and 70 ms, in steps of 20 ms; **Figure S1A**) before and another 12 following adaptation to streams of collision stimuli. Both times, these 192 trials were presented in random order. Test events had a target overlap of 0.25 disc diameters between the first discs final position and the second discs initial position for positive delays whereas for negative delays, the two discs did not overlap. We chose this value to create a strong sense of causal launching at a delay of 0 ms while allowing for the strong effect of adaptation to streams of collision stimuli that we saw in Experiment 1. Adapting streams were shown left and right of fixation for all observers. All other aspects of the procedure and stimuli were identical to those in Experiment 1.

Figure S1C shows the average proportion of causal reports as a function of delay between stopping and starting. Before adaptation (black lines), observers had the strongest impression of causality for slightly positive delays, replicating earlier results [1,14-16]. We captured this relationship with a psychometric function of the form

$$\hat{r}(\delta) = rac{lpha}{2} \left[\operatorname{erf}\left(rac{\delta - \mu_1}{\sqrt{2\sigma^2}}
ight) + \operatorname{erf}\left(rac{\mu_2 - \delta}{\sqrt{2\sigma^2}}
ight)
ight].$$

Here, δ is the delay, α estimates the amplitude of the response function, μ_1 and μ_2 are the means of the two lateral lobes, and σ is their slope. A change of *a* following prolonged exposure to causal collision stimuli would provide additional evidence for the adaptation of causality, independent of a change in perceived timing.

The marked decrease in the average proportion of causal reports, irrespective of the delay between the end of the first disc's motion and the beginning of the second, supports our interpretation of an adaptation of causality rather than of timing. This reduction was spatially specific to the adapted location (**Figure S1C**, blue line, upper panel) and absent at the unadapted location (gray line, lower panel). A bootstrap analysis of the amplitude parameter α in which we resampled each observers individual trials 10,000 times, and re-fitted each of the 10,000 resulting averages [cf. 35] to obtain 95% confidence intervals for α confirmed this result ($\Delta \alpha_{adapted} = -0.312\pm0.128$; $\Delta \alpha_{unadapted}$ $= -0.073\pm0.0836$; $\Delta \alpha_{adapted} - \Delta \alpha_{unadapted} = -0.239\pm0.154$). We conclude, therefore, that the exposure to collisions adapts detectors of causality for this type of stimulus, not detectors of relative timing which in turn feed into a decision regarding the causal nature of events.

We would also like to emphasize that for any timing hypothesis to hold, adaptation to causal launches must affect perceived timing while adaptation to non-causal slips does not. However the timing of the two discs' motion onsets and offsets is the same in both conditions, so a timing hypothesis is plausible only if adaptation is contingent on the spatial configuration that mediates the action transfer—causality—making the adaptation specific to causal interactions rather than temporal factors alone. By showing that localized adaptation occurs for collisions but not for slip events, we isolated a neural substrate for this transfer step that cannot be due to timing, which is equated in the two conditions. Perceptually causal and perceptually non-causal events will always differ in one way or another. We designed the slip event in Experiment 2 such that it was matched in timing but differed in the critical feature: the locations of the trajectory end and start points.



Figure S1 Adapting perceived causality versus perceived timing

(A) Test event. After 80 ms of motion, Disc 1 (black) stopped. With a delay of δ , Disc 2 (blue) moved off in the same direction and at the same speed. Disc 2 started moving either shortly before Disc 1 reached its final position (negative δ) or shortly after (positive δ).

(B) Predictions for the changes in the frequency of causal launch reports following prolonged exposure to streams of collision stimuli. The two markers highlight the data points known from Experiments 1 and 3, a reduction in perceived causality for 0-delay test events from before (black) to after (blue) adaptation. Before adaptation, slightly positive delays would cause the strongest impression of causality [1,14-16]. Following adaptation, two scenarios are compatible with the data from our previous experiments. Adaptation of causality detectors predicts a general decrease in perceived causality (dashed blue line); in contrast, adaptation of perceived timing predicts no decrease of the amplitude of the response function but is compatible with a lateral shift on the time axis and/or a narrowing of the response function (dotted blue line).

(C) Average proportion of causal reports and fits modeling the relationship between timing and perceptual reports.



Figure S2 Distributions of gaze positions during adaptation streams in Experiment 3.

Data are plotted separately for the four observers, accumulated across adaptation streams from all sessions and trials. A dashed red circle surrounds the area covered by stimuli of the adaptation stream; the dashed green circle highlights the fixation area. We counterbalanced adaptation location across observers: MR and AW were adapted to the right of fixation, and NM and AJ were adapted to the left. Contours, derived from kernel density estimation based on linear diffusion processes [36], highlight areas covered by fixation, color-coded by percentage of time spent inside the contour. Each observer's eye position remained inside the fixation area for more than 99% of the time.

2. Supplemental Experimental Procedures

Participants. Eight observers at the University of Potsdam participated in Experiment 1 (27–38 y; mean: 32.3 y; all male) and another eight in Experiment 2 (26–33 y; mean: 30.6 y; 3 female); four participants of these, including two authors, performed both experiments. Each observer ran a one-hour session (432 trials and preceded by 18 warm-up trials). Four observers (20–33 y; mean: 26.5 y; all male; 1 author) at New York University participated in five one-hour sessions (each 288 trials and preceded by 18 warm-up trials) of Experiment 3. All observers had normal or corrected-to-normal vision. Experiments were performed in accordance with the ethical standards laid down in the 1964 Declaration of Helsinki. Potential participants were informed of their right to abstain from participation in the study or to withdraw consent to participate at any time without reprisal. In agreement with the ethics and safety guidelines at the University of Potsdam, we obtained a verbal informed consent statement from all individuals prior to their participation in the study. Participants at New York University signed an informed consent form before study participation.

Setup. Observers sat in a sound-shielded, dimly lit room and received written instruction together with exemplary movies of 0% and 100% overlap events (giving rise to the perception of a causal launch vs. a non-causal pass, respectively). 18 practice trials familiarized them with the procedure. We presented stimuli on a 22" CRT screen (1280x960 pixels, 100 Hz vertical refresh rate; iiyama Vision Master Pro 514 and ~50 cm distance in Experiment 1 and 2, Sony GDM-F520 and 57 cm distance in Experiment 3). In Experiment 3, observers positioned their head on a chin rest and we recorded their right eye's gaze position using an EyeLink 1000 Desktop Mount (SR Research, Osgoode, Ontario, Canada), calibrated before each block of trials. Standard MATLAB (MathWorks, Natick, Massachusetts, USA) toolboxes [37-39] controlled stimulus presentation and response collection by a standard keyboard.

Psychometric function fitting. To determine PSEs, we fitted cumulative Gaussians functions with four parameters (mean μ , standard deviation σ , lower and upper asymptote, γ and λ , respectively) to each observer's reports using maximum-likelihood estimation [40] with no prior assumptions about μ and σ , and with the mean and standard deviation of Gaussian priors for γ and λ set to 0 and 0.05, respectively. The PSE is the magnitude of overlap at which the fit estimates both response alternatives ('launch' vs. 'pass') to be equally likely.

Statistics. Confidence intervals and standard errors (s.e.m.) were bootstrapped using standard procedures. On each bootstrap repetition we resampled *N* PSEs from the *N* observers (with replacement) in a given condition, and then averaged across the 8 resampled PSEs to create the mean of the bootstrap sample. Then, after 10,000 independent repetitions, we computed s.e.m. and 95% confidence interval (1.96 * s.e.m.) from the bootstrap samples [41]. To determine the change in PSE from before to after adaptation at a given location, we computed 95% confidence intervals from 10,000 bootstraps of the observers' Δ PSE at that location. To determine *p*-values for the null hypothesis that the mean difference was not different from 0, we determined what fraction of the bootstrapped distribution lay beyond 0 (assuming normality). To determine the location specificity of a change in PSE, we applied the same procedure to the differences in Δ PSEs between the two test locations, i.e., Δ PSE_{specific} = Δ PSE_{adapted} - Δ PSE_{unadapted}. To compare Δ PSE_{specific} across experiments, we subtracted the 10,000 bootstrap means of Δ PSE_{specific} for Experiments

1 and 2 and based the statistics on the resulting distribution. A geometrical interpretation of this three-way interaction is the distance between the average data points in **Figures 2C** and **3C** along the axis of the diagonal error bars.

We computed a Bayes-Factor (BF) for each reported effect, describing the odds of a difference (no prior assumptions, implemented as a uniform distribution in the range of [-1,1]) relative to the null-hypothesis, given the observed data [42]. Values smaller than 1 favor the absence of an effect; values greater than 1 favor its presence. The likelihoods that we estimated from the data were corrected for small sample sizes; their standard deviations were the estimated s.e.m. multiplied by a factor of $1+20/N^2$, where N is the number of subjects in a given experiment [43].

Eye movement analyses. In Experiment 3 (reference frame experiment), eye position was available during the experiment and correct fixation initiated a trial. Saccade detection and eye position analyses were done offline. Saccades were detected based on their 2D velocity [44], computing smoothed eye velocities using a moving average over five subsequent eye position samples in a trial. Saccades exceeded the median velocity by 5 SDs for at least 8 ms (overshoots were merged with the main saccades if separated by \leq 10 ms). After each adaptation stream, two refixation saccades were required, each within 750 ms of a fixation step, first away from the spot fixated during adaptation, then to the position fixated during the test event. We defined refixation saccades as the first saccades that left the fixation region and landed inside the target area (radii of 2.5° centered on each fixation spot). We rejected trials in which observers blinked or made saccades greater than 1° (in addition to the two refixation saccades) between the first fixation step and the end of the test event. Final data analyses included a total of 5,293 trials (or 91.9%).

Movie S1. Examples of test events, slowed down by a factor of ~3.3

While the observer fixates the center of the screen, test stimuli are presented to the left and to the right of fixation. Here, we show three test events with different amounts of overlap (0.125, 0.5, and 0.875 disc diameters). Note that when there is little overlap (0.125), one disc appears to launch the other's motion; whereas when there is large overlap (0.875), one disc appears to pass over (or under) the other which in turn remains stationary. A moderate amount of overlap (0.5) can evoke either percept. Note that the visual impression of passing is stronger at the original speed of the events.

Movie S2. Collision adaptation stimuli, slowed down by a factor of ~3.3

While the observer fixates the center of the screen, adaptation stimuli are presented to the left and to the right of fixation. Here, we show three short streams at different directions each consisting of two back and forth collisions.

Movie S3. Slip adaptation stimuli, slowed down by a factor of ~3.3

While the observer fixates the center of the screen, adaptation stimuli are presented to the left and to the right of fixation. Here, we show three short streams at different directions each consisting of two back and forth slip stimuli. The motions may appear disorganized but the motion paths are the same as for the collision stimuli except for different end and start locations.

3. Supplemental References

- 35. Herrmann, K., Montaser-Kouhsari, Carrasco, M., and Heeger, D.J. (2010). When size matters: attention affects performance by contrast or response gain. Nat. Neurosci. *13*, 1554-1559.
- 36. Botev, Z.I., Grotowski, J.F., and Kroese, D.P. (2010). Kernel density estimation via diffusion. Ann. Statist. *38*, 2916-2957.
- 37. Brainard, D.H. (1997). The Psychophysics Toolbox. Spat. Vis. 10, 433-436.
- 38. Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat. Vis. *10*, 437-442.
- 39. Cornelissen, F.W., Peters, E.M., and Palmer, J. (2002). The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. Behav. Res. Meth. Instrum. Comput. *34*, 613-617.
- 40. Wichmann, F.A., and Hill, N.J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. Percept. Psychophys. *63*, 1293-1313.
- 41. Efron, B. and Tibshirani, R.J. (1993). An Introduction to the Bootstrap (Chapman and Hall/ CRC).
- 42. Jeffreys, H. (1961). The Theory of Probability (Oxford: Oxford University Press).
- 43. Berry, D.A. (1996). Statistics: A Bayesian Perspective (Duxbury Press).
- 44. Engbert, R., and Mergenthaler, K. (2006). Microsaccades are triggered by low retinal image slip. Proc. Natl. Acad. Sci. USA *103*, 7192-7197.