Visual routines for detecting causal interactions are tuned to motion direction

Sven Ohl[1] & Martin Rolfs[1,2]

[1]Department of Psychology, and [2]Berlin School of Mind and Brain

Humboldt-Universität zu Berlin

Rudower Chaussee 18, 12489 Berlin, Germany

Email:

sven.ohl@hu-berlin.de

martin.rolfs@hu-berlin.de

Running title: Feature-selective causality adaptation

Acknowledgments

This research was supported by a DFG research grant to S.O. (OH 274/4-1) as well as funding from the Heisenberg Programme of the DFG to S.O. (OH 274/5-1) and M.R. (grants RO3579/8-1 and RO3579/12-1). We acknowledge support by the Open Access Publication Fund of Humboldt-Universität zu Berlin. The authors declare no competing financial interests.

Corresponding author:

Sven Ohl
Phone: ++49 (0)30 2093 6789
Email: sven.ohl@hu-berlin.de

Abstract

Detecting causal relations structures our perception of events in the world. Here, we determined for visual interactions whether generalized (i.e., feature-invariant) or specialized (i.e., feature-selective) visual routines underlie the perception of causality. To this end, we applied a visual adaptation protocol to assess the adaptability of specific features in classical launching events of simple geometric shapes. We asked observers to report whether they observed a launch or a pass in ambiguous test events (i.e., the overlap between two discs varied from trial to trial). After prolonged exposure to causal launch events (the adaptor) defined by a particular set of features (i.e., a particular motion direction, motion speed, or feature conjunction), observers were less likely to see causal launches in subsequent ambiguous test events than before adaptation. Crucially, adaptation was contingent on the causal impression in launches as demonstrated by a lack of adaptation in non-causal control events. We assessed whether this negative aftereffect transfers to test events with a new set of feature values that were not presented during adaptation. Processing in specialized (as opposed to generalized) visual routines predicts that the transfer of visual adaptation depends on the feature-similarity of the adaptor and the test event. We show that the negative aftereffects do not transfer to unadapted launch directions but do transfer to launch events of different speed. Finally, we used colored discs to assign distinct feature-based identities to the launching and the launched stimulus. We found that the adaptation transferred across colors if the test event had the same motion direction as the adaptor. In summary, visual adaptation allowed us to carve out a visual feature space underlying the perception of causality and revealed specialized visual routines that are tuned to a launch's motion direction.

Significance statement

We used visual adaptation to carve out the visual feature space critical for detecting collisions in visual launching events. Observers were less likely to report perceiving a collision after the repeated viewing of one disc colliding with and launching the movement of another disc. Importantly, observers' perception of collisions in the direction opposite to the adaptor remained unaffected. However, provided the test stimulus had the adapted direction of motion, the speed or the color of the objects involved in the launching event did not need to be the same as during adaptation for a negative aftereffect to occur. Thus, visual routines that detect causal interactions are selective for motion direction, providing evidence that the perception of causality in a visual interaction relies on low-level perceptual processes.

Keywords

Causality; perception; visual cognition; adaptation; direction selectivity

Introduction

When two objects collide, the objects' future path can be inferred based on physical laws, allowing us, for instance, to purposely change the motion path of an asteroid by steering an uncrewed spacecraft into a collision with the asteroid (see NASA's Double Asteroid Redirection Test; **Figure 1a**). In contrast, how humans detect such causal interactions in their visual environment is less clear. The proposed theories regarding our understanding of causal sensory interactions vary considerably, proposing purely perceptual (Michotte, 1963; Scholl & Tremoulet, 2000), abstract probabilistic (Sanborn et al., 2013), or cognitive mechanisms (Rips, 2011; Weir, 1978; White, 2006). These are neither mutually exclusive nor do they make distinct predictions and they are therefore hard to distinguish (Rips, 2011). Here, we will study the mechanisms underlying the computations of causal interactions in the visual domain by capitalizing on visual adaptation of causality (Kominsky & Scholl, 2020; Rolfs et al., 2013). Adaptation is a powerful behavioral tool for discovering and dissecting a visual mechanism (Kohn, 2007; Webster, 2015) that provides an intriguing testing ground for the perceptual roots of causality. Perceptual accounts of causal understanding posit the existence of visual routines[1] (e.g., cause detectors) located in the visual system, however, their precise nature remain rather hazy. In its purest form such a visual routine could constitute a generalized mechanism that responds to all kinds of causal visual interactions. Alternatively, there could be many different specialized visual routines for the detection of causality that are tuned to the features of a particular causal interaction (e.g., the direction, kinematics, or object identity; **Figure 1b**; see Rips, 2011 for a discussion). Here, we will use visual adaptation to determine whether the computation of causality occurs in generalized (i.e., feature-invariant) or in specialized (i.e., feature-selective) visual routines. Specifically, we will determine whether an adaptor is most effective for test events that match the features of the adaptor.

To this end, we took advantage of the phenomenon that observers perceive causality even in simple kinematic displays—a moving disc that stops next to another disc appears to launch the second disc into motion (i.e., launching effect, **Figure 1c**; Michotte, 1963; see Rips, 2011; Scholl & Tremoulet, 2000; White, 2006, 2017 for reviews). The prolonged viewing of such launching events in an adaptation protocol strongly alters the perception of causality by reducing the proportion of reported launches in subsequent test events (Rolfs et al., 2013). The adaptation of causality is spatially specific to the retinotopic coordinates of the adapting stimulus (Kominsky & Scholl, 2020; Rolfs et al., 2013; for an object-centered elasiticity aftereffect using a related stimulus on a circular motion path, see Arnold et al., 2015), suggesting that the detection of causal interactions is implemented locally in visual space. Here, we selected adaptors from a feature space—while keeping their spatial locations constant—and determined whether adaptation transfers to test events with a different set of features in that feature space (**Figure 1b**). If the strength of the adaptation's perceptual consequences depends on the similarity of the adapter and the test event, it would reveal the existence of specialized visual routines for the detection of causal interactions that are tuned to that feature (i.e., there would be multiple visual routines within a feature dimension, each tuned to a different preferred feature value). Alternatively, if we observe adaptation transfer

[1] We use the term visual routine here to highlight that our adaptation experiments can reveal a causality detection mechanism that resides in the visual system. At the same time, calling it a routine emphasizes similarities with a local, semi-independent operation (e.g., the recognition of familiar motion patterns; see also Ullman, 1987; Cavanagh, Labianca, & Thornton, 2001) that can engage mid- and higher-level processes (e.g., during causal capture, Scholl & Nakayama, 2002; or multisensory integration, Körding et al., 2007).

from one feature to another, this will support the notion of a generalized, feature-invariant, visual routine.

In three experiments, we assessed whether adaptation to launches of a particular motion direction, motion speed, or feature identity (i.e., a conjunction of two features) transfers to other values of that feature space or, alternatively, whether the consequences of adaptation are restricted to the feature values of the adaptor. Our results provide compelling evidence for visual routines that are specialized for processing launches of a particular motion direction.

Method

Participants. In each of the three experiments, we tested eight human observers (**Experiment 1**: ages 21–32 years; 5 female; 3 male; 8 right-handed; 5 right-eye dominant: **Experiment 2**: ages 21–30 years; 5 female; 3 male; 7 right-handed; 6 right-eye dominant; **Experiment 3**: ages 23–30 years; 6 female; 1 non-binary; 7 right-handed; 6 right-eye dominant) in five sessions (1 training session without adaptation, 4 test sessions with adaptation). We determined the final sample size by computing 90 % power contours as a function of sample size and trial number (Baker et al., 2021) and defining the additional criterion of a minimum of 8 observers per experiment. Based on the data obtained in the session without adaptation (i.e., the first session), we determined whether participants do distinguish between passes and launches in the basic experiment by determining whether the proportion of reported passes increases with increasing disc overlap. We excluded observers that did not distinguish between passes and launches after the first session and replaced them by new observers (resulting in the replacement of one observer in each experiment). Data obtained in the first session (i.e., without adaptation) did not enter the final analyses. In **Experiment 1** and **2**, we paid observers 8€ per session as compensation for participation and a bonus of 4€ after successful completion of all sessions. In **Experiment 3**, we paid observers 10€ per session. In all experiments, we obtained observers' written informed consent before the first session. All observers had normal or corrected-to-normal vision and color vision as assessed using the color test *Tafeln zur Prüfung des Farbensinnes/Farbensehens* (Kuchenbecker & Broschmann, 2016). The study was approved by the ethics committee of the Psychology Department of the Humboldt-Universität zu Berlin and it followed the guidelines of the Declaration of Helsinki (2008).

Material. Observers sat in a sound-shielded, dimly lit room putting their head on a chin and forehead rest. We controlled for observers' eye position by tracking their dominant eye using an Eyelink 1000 Desktop Mount eye tracker (SR Research, Ottawa, ON, Canada) with a sampling rate of 1000 Hz. We displayed visual stimuli on a video-projection screen (Celexon HomeCinema, Tharston, Norwich, UK) using a PROPixx DLP projector (VPixx Technologies Inc., Saint Bruno, QC, Canada) at a spatial resolution of 1920 x 1080 pixels and a refresh rate of 120 Hz. The screen was mounted on a wall at 180 cm away from the observer. The experiment was running on a DELL Precision T7810 (Debian GNU Linux 8) and implemented in Matlab 2016b (Mathworks, Natick, MA, USA) using the Psychophysics toolbox 3 (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997) for stimulus presentation and the Eyelink toolbox (Cornelissen et al., 2002) for control of the eye tracker. Behavioral responses were collected by pressing one of the two possible keys on a standard keyboard.

General procedure. In our experiments, we asked observers to report whether they perceived a launch or a pass in test events composed of two discs. In the test events, a peripheral disc

approached a stationary second disc and stopped at varying disc overlaps across trials (**Figure 1c**). Immediately after that the second disc started moving in the same direction and with the same speed as the first disc. Before the first trial, observers read the instruction for the experiment that was displayed on the screen. As part of the instructions, we presented short demo trials of test events with 0 % overlap that are typically perceived as launches (see Supplementary Material Video 1), and test events with full disc overlap, that are typically perceived as a pass (see Supplementary Material Video 2). Observers had the opportunity to inspect these two events as often as they wanted. Following the instructions, observers ran a short training session of 14 trials with varying disc overlaps ranging from 0–100 % overlap.

At the beginning of a trial, we asked observers to fixate a gray fixation point (diameter of 1.5 dva) in the center of the screen on a black background and the trial only started after observers successfully fixated the fixation symbol for at least 200 ms. We presented the test events at 3 dva below the fixation symbol. In these test events, the first moving disc (gray; diameter of 1.5 dva) located either left or right from the vertical meridian started moving towards a stationary disc (gray; diameter of 1.5 dva) located 3 dva below the fixation point. The first disc stopped moving at one of seven possible distances away from the stationary disc, resulting in seven different disc overlaps ranging from 0 to 100 %. The entire duration of the test event was 175 ms. At the end of the trial, observers reported whether they perceived a launch or a pass in the test event by pressing either the arrow up key for launches or arrow down key for passes. In our study, we assessed causal perception by asking observers to report whether they observed a launch or a pass in events of varying ambiguity. This method assumes that launches and passes can be mapped onto a dimension that ranges from causal to non-causal impressions. It has been questioned whether pass events are a natural representative of non-causal events: Observers often report high impressions of causality upon first exposure to pass events, which then decreased after seeing a canonical launch (Bechlivanidis, Schlottmann, & Lagnado, 2019). In our study, therefore, participants completed a separate session that included canonical launches before starting the main experiment.

In addition, the first session served the purpose to determine whether observers perceive launches and passes as a function of disc overlap in our behavioral task. To this end, observers completed 10 blocks with the direction of the test event (2 conditions: left vs. right) as well as the amount of disc overlap (7 conditions: ranging from 0–100 % overlap) presented in a randomized order in a block. Each combination of these manipulations was repeated two times in a block, resulting in 28 trials in each block and a total of 280 trials in the first session. In sessions 2–5 we presented additional adaptors that varied between the experiments.

In all experiments, we tracked observers' dominant eye to ensure proper fixation behavior during presentation of the test events and presentation of the adaptors. More specifically, we tracked the dominant eye's current position at a sampling rate of 1000 Hz and determined online the eyes' distance to the screen center. We aborted a trial, whenever the distance between eye position and screen center exceeded 2 dva. Observers repeated these trials at the end of a block in randomized order. During presentation of the adaptors, we presented a short message (at the fixation point) asking observers to please fixate in the center of the screen once observers gaze exceeded 2 dva away from the screen center.

Adaptation in Experiment 1. In **Experiment 1**, each observer ran in 15 blocks in each of sessions 2–5. The first 5 blocks in a session were again without adaptation to measure an observer's perception of causality before adaptation. As in the first session, we presented test events in two possible horizontal directions and varied the disc overlap resulting in 28 trials in

each block. In blocks 6–15, we presented an adaptor before the first trial of a block. The adaptor was chosen from one of four possible adaptors. In two of the sessions, we used launches in a particular direction as the adaptor (the adaptor's direction was fixed in a session). Before the first trial of a block, we presented 320 launching events at the same location as the test events described above. The exact direction of these launching events was randomly chosen from a narrow uniform distribution around the direction on the horizontal meridian (±30 degrees). The long adaptation phase was complemented by a top-up adaptation of 16 launching events in the same direction as the adaptor (again drawn from the same narrow uniform distribution around that main direction) before each trial to maintain an effective adaptation across the entire block.

In two additional sessions, we presented slip events as adaptors to control that the adaptation was specific for the impression of causality in the launching events (see Supplementary Material Video 3). Slip events are designed to match the launching events in as many physical properties as possible while producing a very different, non-causal phenomenology. In slip events, the first peripheral disc also moves towards a stationary disc. In contrast to launching events, however, the first disc passes the stationary disc and stops only when it is adjacent to the opposite edge of the stationary disc. While slip events do not elicit a causal impression, they have the same number of objects and motion onsets, the same motion direction and speed, as well as the same spatial area of the event as launches. As for the launch-adaptor, we displayed slip adaptors in a narrow uniform range of directions around one of the two possible horizontal directions (from left-to-right or from right-to-left being fixed in a session). This experimental design resulted in four sessions and the order of the type of adaptor (launches vs. slip adaptors; movement direction of the adaptor) was randomly determined for each participant. Overall, each observer ran in a total of 1,680 trials in **Experiment 1**.

Adaptation variations in Experiment 2. In **Experiment 2**, we used a very similar design as in **Experiment 1** with one notable difference. Adaptors and test events varied in motion speed either being the same as in **Experiment 1** or half that speed. Correspondingly, in slow events the test event duration was twice as long (i.e., event duration of 350 ms) as for test events displayed at fast motion speed (i.e., event duration of 175 ms). We presented all test events (and adaptors) always in the same direction (from left to right). Similarly, we also displayed slip events as adaptors in slow and fast motion speeds. In sessions with adaptation, the first five blocks were without adaptation while blocks 6–15 were with adaptors, each block consisting of 28 trials. Each observer ran in a total of 1,680 trials in **Experiment 2**.

Adaptation variations in Experiment 3. In **Experiment 3**, we determined the transfer of adaptation for feature conjunctions (motion direction x color). The two discs in the test events and during adaptation had different colors (i.e., red and green) and could move in both horizontal directions. The red and green discs were not matched for luminance. Measurements obtained after the experiments yielded a luminance of 21 cd/m^2 for the green disc and 6 cd/m^2 for the red disc. In contrast to **Experiment 1 and 2**, slip events were not used as adaptors. This design resulted in four different adaptors that were presented in separate sessions. Again, the first five blocks of a session were without adaptation while blocks 6–15 were with adaptors, each block consisting of 28 trials. Each observer ran in a total of 1,680 trials in **Experiment 3**.

Data analysis. For the statistical analyses and estimation of the psychometric functions we used the *quickpsy* package (Linares & López-Moliner, 2016) and the R environment (R Core Team, 2022). We related disc overlap to the proportion of reported launches using logistic functions with four parameters for the intercept, slope, as well as upper and lower asymptotes. We fitted these functions separately for each observer and condition and obtained the points of subjective equality (PSE; the amount of overlap between the two discs in the test events that result in the same proportion of reported launches and passes). For model fitting, we constrained the range of possible estimates for each parameter of the logistic model. The lower asymptote for the proportion of reported launches was constrained to be in the range 0–0.75, and the upper asymptote in the range 0.25–1. The intercept of the logistic model was constrained to be in the range 1–15, and the slope was constrained to be in the range –20 to –1.

For inferential statistics, we analyzed PSEs using repeated-measures analyses of variance (rmANOVA). Error bars indicate ±1 within-subject standard error of the mean (SEM; Baguley, 2011; Morey, 2008). A significant interaction in the rmANOVA was complemented by running post-hoc paired t-tests. Please note, one observer in **Experiment 3** showed such a strong negative aftereffect when test event and adaptor were identical, such that the upper asymptote was below a proportion of 0.5 (i.e., the point when launches and passes are equally often reported). Instead of a non-identifiable PSE for that observer in that particular condition, we computed inflection points of the psychometric functions in all conditions for that observer which then entered the statistical analyses. Neither of the studies reported in this article was preregistered. The data and analysis code has been deposited at the Open Science Framework and is publicly available https://osf.io/x947m/.

Results

Adaptation of causality is selective for motion direction.

In **Experiment 1**, we tested whether causality is computed separately within visual routines specialized for different motion directions. To this end, we presented observers with brief test events in one of two possible horizontal directions (from left to right, or right to left) in which a peripheral disc moved swiftly towards a stationary one. The first disc stopped moving as soon as the two discs partially overlapped (the amount of overlap was manipulated in seven steps from zero to full overlap, **Figure 1c**). Simultaneously, the second disc started to move along the same trajectory, and we asked observers to report whether they perceived that the first

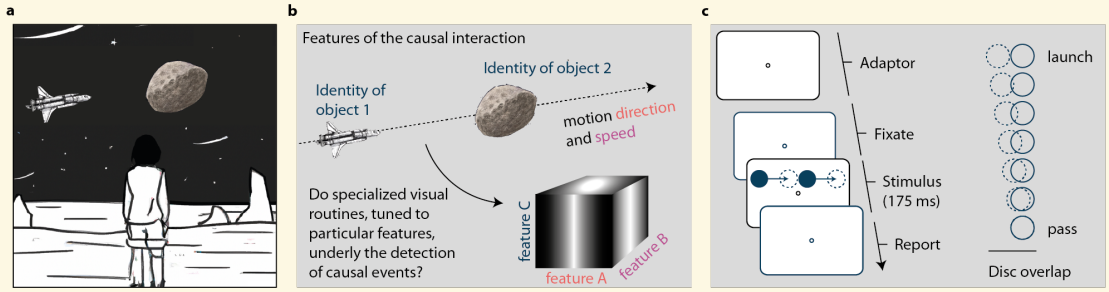

Figure 1. The perception of causality. **a** A person perceiving an upcoming causal interaction between an uncrewed spacecraft and an asteroid. **b** Features of such launching events are the direction and speed of the colliding objects as well as their object identities (i.e., a specific set of features). Assessing the adaptation's transfer between features allows us to determine whether the perception of causality arises in specialized visual routines that are tuned to a particular visual feature. **c** Trial sequence of a test event. A peripheral disc moved towards a stationary disc and stopped with some degree of overlap (ranging from 0–100% overlap in seven equidistant steps) between the two discs. The second disc then immediately started to move in the same direction, with the same speed as the first disc. In adaptation blocks, 320 launches were presented before the first test event of a block, and 16 top-up adaptation events before each subsequent test event.

disc passed the second one (i.e., a pass event; commonly reported at full overlap), or rather launched it into motion (i.e., a launch event; common at zero overlap). We quantified the point of subjective equality (PSE) between perceived launches and passes by modelling psychometric functions for the perception of causality both *before* and *after* adaptation to causal launches (note that we presented the adaptor in a range of ±30° around one of the two possible directions). Moreover, we compared the visual adaptation to launches to a (non-causal) control condition in which we presented slip events as adaptor. In a slip event, the initially moving disc passes completely over the stationary disc, stops immediately on the other side, and then the initially stationary disc begins to move in the same direction without delay. Thus, the two movements are presented consecutively without a temporal gap. This stimulus typically produces the impression of two independent (non-causal) movements.

Visual adaptation to launches successfully affected the perception of causality in this task. We observed a strong negative aftereffect when test events matched the direction of the launch-adaptor. That is, observers were less likely to report a launch after adapting to launches in the same direction as the test events (**Figure 2a–c**). This observation was corroborated by a two-way (2 event types: launch vs. slip; 2 directions: congruent vs. incongruent) repeated measures analysis of variance (rmANOVA) in which we assessed the magnitude of adaptation by subtracting the PSE before adaptation (PSE_{before} = 0.60, $CI_{95\%}$ = [0.59 0.62]) from the PSEs obtained in each of the four different adaptation conditions (PSE_{cong_launch} = 0.40, $CI_{95\%}$ = [0.36 0.43]; PSE_{incong_launch} = 0.55, $CI_{95\%}$ = [0.54 0.56]; PSE_{cong_slip} = 0.55, $CI_{95\%}$ = [0.54 0.57]; PSE_{incong_slip} = 0.61, $CI_{95\%}$ = [0.59 0.64]). The analysis revealed a significant main effect of event type (F (1, 7) = 62.73, $p < 0.001$; **Figure 2a–c**) demonstrating stronger adaptation for launching than slip events (ΔPSE_{event} = –0.11, $CI_{95\%}$ = [–0.14 –0.08]). Moreover, adaptors in a congruent direction as the test event resulted in stronger adaptation than in incongruent directions (F (1, 7) = 41.62, $p < 0.001$; $\Delta PSE_{direction}$ = –0.11, $CI_{95\%}$ = [–0.15 –0.07]). Critically, we observed a significant interaction (F (1, 7) = 12.23, p = 0.01), revealing

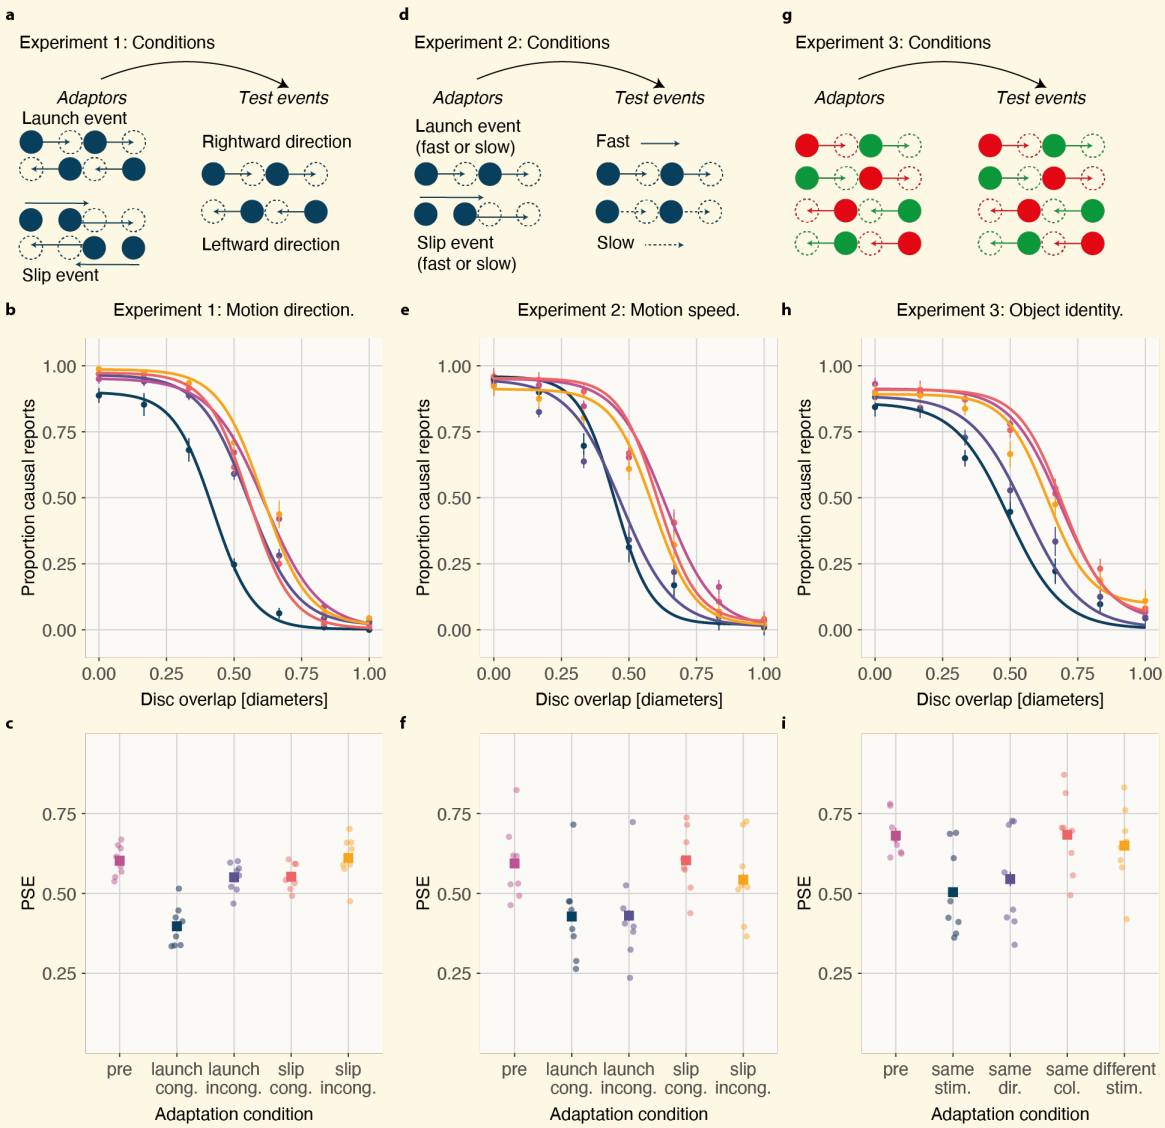

Figure 2. Results of Experiment 1, 2, and 3. **a** Experimental conditions of Experiment 1. **b** Mean proportion of causal reports as a function of disc overlap in Experiment 1. Visualization of psychometric curves is based on the mean parameters averaged across observers before adaptation (pink), and after adaptation with direction-congruent launches (blue), direction-incongruent launches (in purple), direction-congruent slip events (red) and direction-incongruent slip events (in orange). **c** PSEs for each individual observers (circles) and the mean across observers (square). **d** Experimental conditions of Experiment 2. **e** Mean proportion of causal reports as a function of disc overlap in Experiment 2. Visualization of psychometric curves is based on the mean parameters averaged across observers before adaptation (pink), and after adaptation with speed-congruent launches (blue), speed-incongruent launches (in purple), speed-congruent slip events (red) and speed-incongruent slip events (in orange). **f** PSEs for each individual observers (circles) and the mean across observers (square). **g** Experimental conditions of Experiment 3. **h** Mean proportion of causal reports as a function of disc overlap in Experiment 3. Visualization of psychometric curves is based on the mean parameters averaged across observers before adaptation (pink), and after adaptation with adaptors that are identical with the test event (blue), share the same direction but different colors (in purple), same colors but different direction (red) and completely different test events (in orange). **i** PSEs for each individual observers (circles) and the mean across observers (square). All error bars are ± 1 SEM.

that the influence of direction-congruency on the magnitude of adaptation was significantly stronger for launches than for slip events (Post-hoc t-test: $t(7) = 3.50$, $p = 0.01$; $\Delta PSE_{eventXdirection} = -0.09$, $CI_{95\%} = [-0.16\ -0.03]$) supporting the hypothesis that specialized visual routines are tuned to the motion direction of launching events.

Adaptation of causality transfers across motion speed.

In **Experiment 2**, we determined whether the perception of causality is also selective for other feature dimensions that are relevant for causal interactions. Causal events in our visual environment encompass interactions of various motion speeds. Here, we examined whether adapting the perception of causality transfers between events that either had different speeds (same as in **Experiment 1** or half that). Both speed-congruent (i.e., test events had the same speed as the adaptor) and speed-incongruent launches (i.e., test events and adaptor differed by a factor of 2) resulted in adaptation, demonstrating a transfer of adaptation across motion speeds (**Figure 2d–f**). As in **Experiment 1**, we corroborated this observation by a two-way rmANOVA in which we assessed differences in the magnitude of adaptation (i.e., the difference in PSEs before and after adaptation; PSE_{before} = 0.59, $CI_{95\%}$ = [0.57 0.61]; PSE_{cong_launch} = 0.43, $CI_{95\%}$ = [0.40 0.46]; PSE_{incong_launch} = 0.43, $CI_{95\%}$ = [0.39 0.47]; PSE_{cong_slip} = 0.60, $CI_{95\%}$ = [0.57 0.63]; PSE_{incong_slip} = 0.54, $CI_{95\%}$ = [0.50 0.59]) as a function of speed congruency (congruent vs. incongruent speed), event type (launch vs. slip event) and their interaction. The analysis revealed a significant main effect of event type (F (1, 7) = 23.39, p = 0.002; **Figure 2d–f**) demonstrating stronger adaptation for launching than slip events (ΔPSE_{event} = –0.14, $CI_{95\%}$ = [–0.21 –0.07]). Importantly, the magnitude of adaptation was not significantly different for congruent and incongruent speed between adaptor and test event (F (1, 7) = 3.43, p = 0.107; ΔPSE_{speed} = 0.03 $CI_{95\%}$ = [–0.01 0.07]). Moreover, there was no significant interaction between event type and speed (F (1, 7) = 3.48, p = 0.105; $\Delta PSE_{eventXspeed}$ = –0.06, $CI_{95\%}$ = [–0.14 0.02]). These findings demonstrate that the perception of causality is not tuned to motion speed—or, at least, that the tuning is very broad.

Adaptation of causality using feature-conjunctions.

In **Experiment 3**, we determined whether the identities of the two objects involved in the causal interaction can break the observed dominance of the direction-specific computation underlying the perception of causality. For instance, if two visually distinct objects are involved in a causal interaction with one type of object always launching a second, visually distinct object, it is unclear whether the influence of the adaptor is confined to events with the same feature conjunction or, alternatively, whether motion direction is the only relevant feature dimension for determining the presence of causal interactions. To distinguish between these alternatives, the two discs in the presented test events had two different colors (i.e., they were either red or green). Moreover, the adaptor in a session displayed only one particular feature conjunction (e.g., the adaptor in one session was always a red disc on the left launching a green disc on the right into rightward motion; adaptors varied across sessions and each subject saw each combination of feature conjunction across the multiple sessions). In contrast to the adaptor's fixed feature conjunction in a session, both motion direction and color identities in the test events varied randomly from trial to trial. Again, we observed strong adaptation when test events and adaptors were identical. However, when one of the features in the feature conjunction differed, we observed adaptation only for test events in the same direction as the adaptor, irrespective of the object's color. Thus, color does not constitute a critical feature of the visual routine for detecting causal interactions (**Figure 2g–i**). Again, we corroborated this observation by a two-way rmANOVA in which we assessed differences in the magnitude of adaptation (i.e., the difference in PSEs before and after adaptation; PSE_{before} = 0.68, $CI_{95\%}$ = [0.63 0.73]; $PSE_{same_stimulus}$ = 0.50, $CI_{95\%}$ = [0.47 0.53]; $PSE_{same_direction}$ = 0.55, $CI_{95\%}$ = [0.50 0.59]; PSE_{same_color} = 0.68, $CI_{95\%}$ = [0.65 0.72]; $PSE_{different_stimulus}$ = 0.65, $CI_{95\%}$ = [0.61 0.69]) as a function of motion direction (same vs. different), color assignment (same vs. different) and their interaction. The analysis revealed a significant main effect of motion

direction (F (1, 7) = 21.07, p = 0.003) demonstrating stronger adaptation for launches in the same as compared to the opposite direction as the adaptor ($\Delta PSE_{direction}$ = –0.14, $CI_{95\%}$ = [–0.22 –0.07]). Adaptation following adaptors with the same color vs. different color assignment were not significantly different (F (1, 7) = 0.071, p = 0.798; ΔPSE_{color} = –0.004, $CI_{95\%}$ = [–0.03 0.03]). However, we observed a marginally significant interaction between color and motion direction (F (1, 7) = 4.59, p = 0.07; $\Delta PSE_{directionXcolor}$ = –0.07, $CI_{95\%}$ = [–0.16 0.01]). Post-hoc paired t-tests showed that color had a marginally significant influence on the magnitude of adaptation when adaptor and test events were in same direction (t (7) = 1.97, p = 0.090; $\Delta PSE_{color_samedir}$ = –0.04, $CI_{95\%}$ = [–0.09 0.01]). However, color had no significant influence on the adaptation of causality when adaptor and test events were in opposite direction (t (7) = 1.97, p = 0.090; $\Delta PSE_{color_opp_dir}$ = 0.03, $CI_{95\%}$ = [–0.02 0.09]).

Discussion

We provide new evidence for the fundamental role of perceptual processes in detecting cause and effect. Using visual adaptation, we revealed that visual routines underlie the perception of causality that are specialized for a launch's motion direction but invariant across a range of motion speeds. The tuning of causal perception to motion direction reveals a mechanism that is operating locally in feature space and complements previous reports of a spatially specific mechanism residing in retinotopic space (Kominsky & Scholl, 2020; Rolfs et al., 2013).

The observed adaptation is specific to launching events that evoked a phenomenological impression of causality. This implies that adapting simply to the direction of a moving stimulus will only change the input to the visual routine but not the functioning of the routine itself. We controlled for such potential non-causal adapation of low-level featues (e.g., number of stimuli, time of contact, speed, overlap, or simply motion) using carefully designed control events (i.e., slip events) that did not result in comparable aftereffects.

Motion direction constitutes a basic computational unit in vision and can be computed as early as in retinal circuits (e.g., in retinal ganglion cells about 2–3 synapses downstream of photoreceptors in rabbits, Barlow & Hill, 1963). While a retinal origin of launch detection is unlikely, the specialization of visual routines for the perception of causality at the level of individual motion directions raises the possibility that this function is located surprisingly early in the visual system as opposed to a higher-level visual computation. A similar adaptation-based rationale can be applied to study whether different causal interactions can be distinguished from each other by assessing the transfer of adaptation between different causal interactions (Kominsky & Scholl, 2020). While Michotte speculated about three separate causal impressions (i.e., launching, entraining, triggering; see White, 2017 for an extended catalogue of causal interactions) adaptation helped to refine these categories as an adaptation to triggering events transferred to launching events, but adaptation to entraining had no such influence (Kominsky & Scholl, 2020). Thus, adaptation at the category-level of the causal interaction allows to distinguish spezialized detectors for a particular causal interaction. Our findings break down this specialization to a more fundamental level, showing that a specialized visual routine for launching events exists even within separate motion direction channels. While the present study demonstrates direction-selectivity for the detection of launches, previous adaptation protocols demonstrated successful adaptation using adaptors with random motion direction (Rolfs et al., 2013; Kominsky & Scholl, 2020). These results therefore suggest independent direction-specific routines, in which adaptation to launches in one

direction does not counteract an adaptation to launches in the opposite direction (as for example in opponent color coding).

Habituation studies in six-months-old infants also demonstrated that the reversal of a launch resulted in a recovery from habituation to launches (while a non-causal control condition of delayed-launches did not; Leslie & Keeble, 1987). In their study, the reversal of motion direction was accompanied by a reversal of the color assignment to the cause-effect-relationship. In contrast, our findings suggest, that in adults color does not play a major role in the detection of a launch. Future studies should further delineate similarities and differences obtained from adaptation studies in adults and habituation studies in children (e.g., Kominsky et al., 2017; Kominsky & Scholl, 2020).

The adaptation transferred across different speeds and colors as long as the test events shared the same motion direction as the adaptor. For instance, we demonstrated a transfer of adaptation across speed with symmetrical speed ratios. This result complements a previous finding that reported that the adaptation to triggering events (with a speed ratio of 1:3) resulted in significant retinotopic adaptation of ambiguous (launching) test events of different speed ratios (i.e., test events with a speed ratio of 1:1 and of 1:3; Kominsky & Scholl, 2020). Moreover, perceiving a red disc that launched a green disc into motion is also affected by the adaptation of a green disc launching a red disc (i.e., the opposite color assignment) into motion as long as the two events shared the same motion direction. This finding rules out an alternative cognitive mechanism, in which the causing-disc is first identified (e.g., the red disc is identified to launch other discs into motion) and the criterion for detecting a causal interaction would be selectively adjusted only for this particular causing-disc in a top-down manner. Indeed, our results show that the appearance of the objects in the launching event plays only a minor role for adapting the perception of causality (if any), suggesting that the visual system can calibrate the detection of causal interactions independent of the involved object identities. Our findings therefore provide additional support for the claim that an event's spatiotemporal parameters mediate the perception of causality (Michotte, 1963; Leslie, 1984; Scholl & Tremoulet, 2000).

We suggest that at least two functional benefits result from a specialized visual routine for detecting causality. First, a direction-selective detection of launches allows adaptation to occur separately for each direction. That means that the visual system can automatically calibrate the sensitivity of these visual routines in response to real-world statistics. For instance, while falling objects drop vertically towards the ground, causal relations such as launches are common in horizontal directions moving along a stable ground. Second, we think that causal visual events are action-relevant, and the faster we can detect such causal interactions, the faster we can react to them. Direction-selective motion signals are available very early on in the visual system. Visual routines that are based on these direction-selective motion signals may enable faster detection. While our present findings demonstrate direction-selectivity, they do not pinpoint where exactly that visual routine is located. It is possible that the visual routine is located higher up in the visual system (or distributed across multiple levels), relying on a direction-selective population response as input.

The existence of specialized visual routines computing causality, however, does not rule out the possibility of a more complex hierarchical architecture that is composed of both specialized and generalized routines. The output of local, specialized detectors within that architecture could feed into a global, generalized detector. Thus, a weaker response from adapted specialized cause detectors would elicit a weaker response in unadapted global detectors (or super-detectors; Rips, 2011). Indeed, visual adapation to other low-level visual

features (e.g., contrast sensitivity) can yield weaker responses during early visual processing which are then passed on to other downstream areas (Kohn, 2007).

We used visual adaptation to carve out a bottom-up visual routine for detecting causal interactions in form of launching events. However, we know that more complex behaviors of perceiving causal relations can result from integrating information across space (e.g., in causal capture; Scholl & Nakayama, 2002), across time (postdictive influence; Choi & Scholl, 2006), and across sensory modalities (Sekuler, Sekuler, & Lau, 1997). Bayesian causal inference has been particularly successful as a normative framework to account for multisensory integration (Körding et al., 2007; Shams & Beierholm, 2022). In that framework, the evidence for a common-cause hypothesis is competing with the evidence for an independent-causes hypothesis (Shams & Beierholm, 2022). The task in our experiments could be similarly formulated as two competing hypotheses for the second disc's movement (i.e., the movement was caused by the first disc vs. the second disc did not move). This framework also emphasizes the distributed nature of the neural implementation for solving such inferences, showing the contributions of parietal and frontal areas in addition to sensory processing (for review see Shams & Beierholm, 2022). Moreover, even visual adaptation to contrast in mouse primary visual cortex is influenced by top-down factors such as behavioral relevance—suggesting a complex implementation of the observed adaptation results (Keller et al. 2017). The present experiments, however, presented purely visual events that do not require an integration across processing domains. Thus, the outcome of our suggested visual routine can provide initial evidence from within the visual system for a causal relation in the environment that may then be integrated with signals from other domains (e.g., auditory signals). Determining exactly how the perception of causality relates to mechanisms of causal inference and the neural implementation thereof is an exciting avenue for future research. Note, however, that perceived causality can be distinguished from judged causality: Even when participants are aware that a third variable (e.g., a color change) is the best predictor of the movement of the second disc in launching events, they still perceive the first disc as causing the movement of the second disc (Schlottmann & Shanks, 1992).

Neurophysiological studies support the view of distributed neural processing underlying sensory causal interactions with the visual system playing a major role. Imaging studies in particular revealed a network for the perception of causality that is also involved in action observation (Blakemore et al., 2003; Fonlupt, 2003; Fugelsang et al., 2005; Roser et al., 2005). The fact that visual adaptation of causality occurs in a retinotopic reference frame emphazises the role of retinotopically organized areas within that network (e.g., V5 and the superior temporal sulcus). Interestingly, single cell recordings in area F5 of the primate brain revealed that motor areas are contributing to the perception of causality (Caggiano et al., 2016; Rolfs, 2016), emphasizing the distributed nature of the computations underlying causal interactions. This finding also stresses that the detection, and the prediction, of causality is essential for processes outside purely sensory systems (e.g., for understanding other's actions, for navigating, and for avoiding collisions). The neurophysiology subserving in causal inference further extend the candidate cortical areas that might contibute to the detection of causal relations, emphasizing the role of the frontal cortex for the flexible integration of multisensory representations (Cao et al., 2019; Coen et al., 2023).

Visual adaptation of causality constitutes a powerful tool that allowed us to reveal how specialized visual routines are tuned to the motion direction of a launching event. Visual adaptation in general is thought to be one of the signatures of a perceptual process (Hafri &

Firestone, 2021) and therefore is informative for identifiying the visual origin of a mechanism as opposed to, for instance, cognitive accounts (e.g., cognitive schema; see Rips, 2011 for a review). The detection of causal interactions features many hallmarks of perception: An impression of causality emerges fast and automatically for briefly presented stimuli (Michotte, 1963) and causal events reach awareness faster than non causal-events (Moors et al., 2017). The implementation of visual routines for the perception of causality in separate channels of a basic visual feature such as motion direction fortifies the view that our understanding of causal interactions indeed starts as a low-level perceptual routine.

References

Baker, D. H., Vilidaite, G., Lygo, F. A., Smith, A. K., Flack, T. R., Gouws, A. D., & Andrews, T. J. (2021). Power contours: Optimising sample size and precision in experimental psychology and human neuroscience. *Psychological methods*, *26*(3), 295.

Barlow, H. B., & Hill, R. M. (1963). Selective sensitivity to direction of movement in ganglion cells of the rabbit retina. *Science*, *139*(3553), 412–414. https://doi.org/10.1126/science.139.3553.412

Bechlivanidis, C., Schlottmann, A., & Lagnado, D. A. (2019). Causation without realism. *Journal of Experimental Psychology: General*, *148*(5), 785–804.

Blakemore, S.-J., Boyer, P., Pachot-Clouard, M., Meltzoff, A., Segebarth, C., & Decety, J. (2003). The detection of contingency and animacy from simple animations in the human brain. *Cerebral Cortex*, *13*(8), 837–844. https://doi.org/10.1093/cercor/13.8.837

Brainard, D. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Caggiano, V., Fleischer, F., Pomper, J. K., Giese, M. A., & Thier, P. (2016). Mirror neurons in monkey premotor area F5 show tuning for critical features of visual causality perception. *Current Biology*, *26*(22), 3077–3082. https://doi.org/10.1016/j.cub.2016.10.007

Cao, Y., Summerfield, C., Park, H., Giordano, B. L., & Kayser, C. (2019). Causal inference in the multisensory brain. *Neuron*, *102*(5), 1076–1087.

Cavanagh, P., Labianca, A. T., & Thornton, I. M. (2001). Attention-based visual routines: sprites. *Cognition*, *80*(1-2), 47–60.

Choi, H., & Scholl, B. J. (2006). Perceiving causality after the fact: Postdiction in the temporal dynamics of causal perception. *Perception*, *35*(3), 385–399.

Coen, P., Sit, T. P., Wells, M. J., Carandini, M., & Harris, K. D. (2023). Mouse frontal cortex mediates additive multisensory decisions. *Neuron*, *111*(15), 2432–2447.

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink Toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 613–617. https://doi.org/10.3758/BF03195489

Fonlupt, P. (2003). Perception and judgement of physical causality involve different brain structures. *Cognitive Brain Research*, *17*(2), 248–254. https://doi.org/10.1016/S0926-6410(03)00112-5

Fugelsang, J. A., Roser, M. E., Corballis, P. M., Gazzaniga, M. S., & Dunbar, K. N. (2005). Brain mechanisms underlying perceptual causality. *Cognitive Brain Research*, *24*(1), 41–47. https://doi.org/10.1016/j.cogbrainres.2004.12.001

Hafri, A., & Firestone, C. (2021). The perception of relations. *Trends in Cognitive Sciences*, *25*(6), 475–492. https://doi.org/10.1016/j.tics.2021.01.006

Kleiner, M., Brainard, D. H., Pelli, D. G., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, *36*(14), 1–16.

Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology*, *97*(5), 3155–3164. https://doi.org/10.1152/jn.00086.2007

Kominsky, J. F., Strickland, B., Wertz, A. E., Elsner, C., Wynn, K., & Keil, F. C. (2017). Categories and constraints in causal perception. *Psychological Science*, *28*(11), 1649–1662.

Kominsky, J., & Scholl, B. J. (2020). Retinotopic adaptation reveals distinct categories of causal perception. *Cognition*, *203*, 104339. https://doi.org/10.1016/j.cognition.2020.104339

Körding, K. P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., & Shams, L. (2007). Causal inference in multisensory perception. *PLoS one*, *2*(9), e943.

Kuchenbecker, J., & Broschmann, D. (Eds.). (2016). *Tafeln zur Prüfung des Farbensinnes/Farbensehens*. Georg Thieme Verlag.

Leslie, A. M. (1984). Spatiotemporal continuity and the perception of causality in infants. *Perception*, *13*(3), 287–305.

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality?. *Cognition*, *25*(3), 265–288.

Linares, D., & López-Moliner, J. (2016) *quickpsy: An R Package to Fit Psychometric Functions for Multiple Groups.* The R Journal, 8(1), 122-131.

Michotte, A. (1963). *The perception of causality*. Basic Books. (Original published in 1946).
Moors, P., Wagemans, J., & de-Wit, L. (2017). Causal events enter awareness faster than non-causal events. *PeerJ*, *5*, e2932. https://doi.org/10.7717/peerj.2932
Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.
R Core Team (2022). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.
Rips, L. J. (2011). Causation from perception. *Perspectives on Psychological Science*, *6*(1), 77–97.
Rolfs, M. (2016). Visual neuroscience: Seeing causality with the motor system? *Current Biology*, *26*(22), R1183–R1185. https://doi.org/10.1016/j.cub.2016.09.046
Rolfs, M., Dambacher, M., & Cavanagh, P. (2013). Visual adaptation of the perception of causality. *Current Biology*, *23*(3), 250–254. https://doi.org/10.1016/j.cub.2012.12.017
Roser, M. E., Fugelsang, J., Dunbar, K. N., Corballis, P. M., & Gazzaniga, M. (2005). Dissociating processes supporting causal perception and causal inference in the brain. *Neuropsychology*, *19*(5), 591–602. https://doi.org/10.1037/0894-4105.19.5.591
Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2013). Reconciling intuitive physics and Newtonian mechanics for colliding objects. *Psychological Review*, *120*(2), 411–437. https://doi.org/10.1037/a0031912
Schlottmann, A., & Shanks, D. R. (1992). Evidence for a distinction between judged and perceived causality. *The Quarterly Journal of Experimental Psychology*, *44*(2), 321–342.
Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, *4*(8), 299–309. https://doi.org/10.1016/S1364-6613(00)01506-0
Scholl, B. J., & Nakayama, K. (2002). Causal capture: Contextual effects on the perception of collision events. *Psychological Science*, *13*(6), 493–498.
Sekuler, R., Sekuler, A. B., & Lau, R. (1997). Sound changes perception of visual motion. *Nature*, *384*, 308–309.
Shams, L., & Beierholm, U. (2022). Bayesian causal inference: A unifying neuroscience theory. *Neuroscience & Biobehavioral Reviews*, *137*, 104619.
Ullman, S. (1987). Visual routines. In *Readings in computer vision* (pp. 298-328). Morgan Kaufmann.
Webster, M. A. (2015). Visual adaptation. *Annual Review of Vision Science*, *1*(1), 547–567. https://doi.org/10.1146/annurev-vision-082114-035509
Weir, S. (1978). The perception of motion: Michotte revisited. *Perception*, *7*(3), 247–260. https://doi.org/10.1068/p070247
White, P. A. (2006). The causal asymmetry. *Psychological Review*, *113*(1), 132–147. https://doi.org/10.1037/0033-295X.113.1.132
White, P. A. (2017). Visual impressions of causality. In M. R. Waldmann (Hrsg.), *The Oxford Handbook of Causal Reasoning* (Bd. 1). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199399550.013.17