Cognition 279 (2027) 106773



Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit



Full Length Article

The spatial specificity and recovery from visual adaptation in causality perception

Laura van Zantwijk [a,*], Martin Rolfs [a], Sven Ohl [a,b]

[a] Department of Psychology, Humboldt-Universität zu Berlin, Rudower Chaussee 18, 12489 Berlin, Germany
[b] Department of Applied Human Sciences, Magdeburg-Stendal University of Applied Sciences, Osterburger Str. 25, 39576 Stendal, Germany

ARTICLE INFO

Dataset link: Experiments 1, 2 and 3 (Original data and code)

Keywords:
Causality perception
Visual adaptation
Recovery
Spatial specificity

ABSTRACT

When one object approaches another object which, upon touching, moves in the same direction, humans report a vivid impression of one launching the other. Visual adaptation can alter this perception of causality: observers less often report seeing a launch after viewing a stream of launch events. In three experiments, we further characterized how visual adaptation influences the perception of causality by determining the spatial specificity of adaptation and timecourse of recovery from adaptation. In Experiment 1, observers saw ambiguous test events (i.e., the overlap between the two objects varied over trials) at three different distances from the adapted location. Adaptation was strongest when the adaptor and test event were presented at the same location, and absent when the two were separated by just three degrees of visual angle. Moreover, the perception of causality gradually recovered from adaptation, but remained incomplete. In Experiment 2, both long and short adaptation sequences were highly effective in driving adaptation, and showed no difference in the recovery timecourse, which was complete following more experimental blocks. In Experiment 3, a break without any task-relevant visual input also led to a recovery over the same timespan, but this time, the recovery was instantaneous and incomplete. Altogether, our results provide evidence for highly spatially specific computations, instantaneously responding to the onset of adaptation and then recovering from the adaptation over a short time window.

1. Introduction

Perception adapts quickly to continuous features in the visual environment—a process known as visual adaptation, in which prolonged exposure to a stimulus temporarily alters subsequent perception. If we stare at the downward motion of a waterfall, and then look at a stationary surface, it appears to move upwards (hence, the waterfall illusion; Addams, 1834). And when we put on sunglasses, the world initially appears tinted by the lenses' color, but this tint gradually fades over time (La Hire, 1694/2001; Tregillus & Engel, 2019). These effects typically do not last long, as our system is continuously re-calibrating to the current state of the world. There are, of course, rare exceptions to this general pattern. In some cases, such as prolonged fixation on a moving spiral, motion aftereffects can be observed 20 h later when the adapted retinal region is stimulated again (Masland, 1969). Similarly, the McCollough effect can persist for days or even months (Jones & Holding, 1975).

Strikingly, visual adaptation extends even to seemingly more high-level visual features, including causal interactions (Fig. 1a): repeatedly looking at one disc approaching a second disc, setting the second disc into motion, renders later ambiguous test events perceptually less causal (Kominsky & Scholl, 2020; Kominsky & Wenig, 2025; Ohl & Rolfs, 2025; Rolfs, Dambacher, & Cavanagh, 2013). Repeated exposure to one object triggering another's—where the second motion is considerably faster than the first—results in similar negative aftereffects (Kominsky & Scholl, 2020). These results have supported the view that we understand causal interactions automatically through perceptual processes (Michotte, 1963/2017; Scholl & Tremoulet, 2000) rather than by cognitive reasoning (Rips, 2011; Weir, 1978; White, 2006). This poses questions about the specificity of such adaptation effects in space and time. How does our visual system recover from being adapted to causal interactions, and how spatially precise are the effects of visual adaptation to causal interactions to begin with?

In the present study, we build on the notion that perceiving causal interactions is a perceptual, automatic process that originates in the visual system (Scholl & Tremoulet, 2000). Visual adaptation is a powerful behavioral tool for probing the mechanisms of perception (Kohn, 2007; Webster, 2015), which provides a testing ground for the

* Corresponding author.
E-mail addresses: laura.van.zantwijk@hu-berlin.de (L. van Zantwijk), martin.rolfs@hu-berlin.de (M. Rolfs), sven.ohl@h2.de (S. Ohl).

https://doi.org/10.1016/j.cognition.2026.106773
Received 20 April 2026; Received in revised form 27 October 2026; Accepted 29 October 2026
0010-0277/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).
Open access funding enabled and organized by Projekt DEAL.

processes underlying causal perception. More specifically, using visual adaptation we can present a specific adaptor to observers and test whether adaptation transfers to a test event that differs in some aspect from the adaptor. Depending on whether an adaptation transfer is observed, we can then conclude if this particular aspect of the adaptor is critical for the perceptual mechanism of interest or not. Using this approach, visual adaptation to causal events has been shown to operate in a retinotopic reference frame, that is, it is tied to specific locations relative to the fovea (Kominsky & Scholl, 2020; Kominsky & Wenig, 2025; Rolfs et al., 2013). However, these previous studies have primarily tested spatial specificity by comparing adaptation and test locations that were far apart or even in opposite hemifields. Thus, although these findings demonstrate that causal adaptation is spatially selective, they leave the spatial scale over which adaptation operates unresolved: whether it is broadly constrained to a hemifield or whether it is localized more narrowly, within a visual hemifield. Our measurements of the spatial extent of adaptation therefore provide a more fine-grained test of the spatial organization of the mechanisms underlying the perception of causality. Further evidence for an automatic process of perceiving causal interactions comes from the lack of transfer of visual adaptation across motion directions (Ohl & Rolfs, 2025), suggesting that causality perception relies on direction-selective channels in the visual system. Taken together, such findings support the notion that causality is not inferred solely through higher-level reasoning, but might already be decoded at early stages of visual processing. Whether the effects observed from visual adaptation reflect sensory adaptation, predictive processes, decisional recalibration, or some combination thereof remains an open question. Characterizing the spatial and temporal properties of this phenomenon is therefore a necessary step toward distinguishing between these possibilities.

The suggestion that the perception of causality is the result of low-level visual processes raises the possibility that visual routines underlying the perception of causality are highly spatially specific to a small region of the visual field. The size of receptive fields increases for later processing stages in higher order visual areas (Freeman & Simoncelli, 2011; Wandell, Dumoulin, & Brewer, 2007). Consequently, a rather broad spatial tuning could argue against the notion that causality perception is an early visual process. One way to address this question is by examining the size of the adapted field. If neurons involved in detecting causal interactions have large receptive fields, adaptation should generalize across a wider area of the visual field. Alternatively, if the receptive fields of these neurons are small, adaptation should be tightly confined to the spatial location of the adapting stimulus. Examining the size of the adapted field may therefore provide indirect evidence about the level of the visual hierarchy at which processes contributing to causal perception might be implemented.

Another aspect of causal perception that remains poorly understood is the temporal dynamics of its underlying visual mechanisms. In other words, after adapting to causal interactions, does perception completely return to baseline, or do residual effects persist? Answering these questions would reveal what kind of information the visual system tracks and the timescales over which it calibrates (Webster, 2015). Given that both the environment and the internal state of the observer vary across multiple timescales, it is reasonable that adaptation mechanisms should also operate with different temporal profiles (Kording, Tenenbaum, & Shadmehr, 2007). One possibility is that the visual system responds only briefly to causal interactions, producing rapid effects that quickly decay. Such short-lived adaptation would suggest a flexible and responsive mechanism, capable of rapidly adjusting to moment-to-moment changes. Alternatively, adaptation might persist over longer durations, reflecting a more stable recalibration—that is, a sustained shift in the mapping between sensory input and perceptual judgment following recent visual experience, rather than implying that perception has become optimally adjusted to the environment.

An additional open question concerns the timecourse of the recovery from adaptation: does the visual system return to baseline immediately, in a step-like fashion, or does it recover slowly over time? While there are many findings on the onset of various visual adaptation effects, there is far less research on how the system recovers from them. One exception is the phenomenon of color afterimages. After fixating on a colored pattern, observers initially perceive a complementary-colored afterimage when the stimulus is removed, but this percept gradually fades and the visual system returns to its baseline state (Zaidi, Ennis, Cao, & Lee, 2012). In motor learning, recovery dynamics have been instrumental for modeling adaptive processes, prompting the development of targeted protocols within the oculomotor domain (Cassanello, Ohl, & Rolfs, 2016). Despite this, the timecourse of the recovery from visual adaptation remains largely understudied. Understanding the recovery profile can give insights into how enduring the visual aftereffects of causal adaptation are, and whether they reflect short-term changes in perceptual recalibration or longer-lasting perceptual shifts.

The current experiments had three objectives. First, we assessed how far away from the spatial location of visual adaptation the perception of causal interactions is still affected. Second, we asked how strong an adaptor we need to observe a shift in the perception of causality, and how quickly this effect might recover back to the starting point. Finally, we asked whether recovery from visual adaptation occurs in a task-relevant manner, through active recalibration of the visual system to the statistical regularities of the environment, or in a time-based manner, resembling a more passive form of recalibration.

To realize this, we made use of the phenomenon that observers readily perceive causality in simple kinematic displays; a moving disc

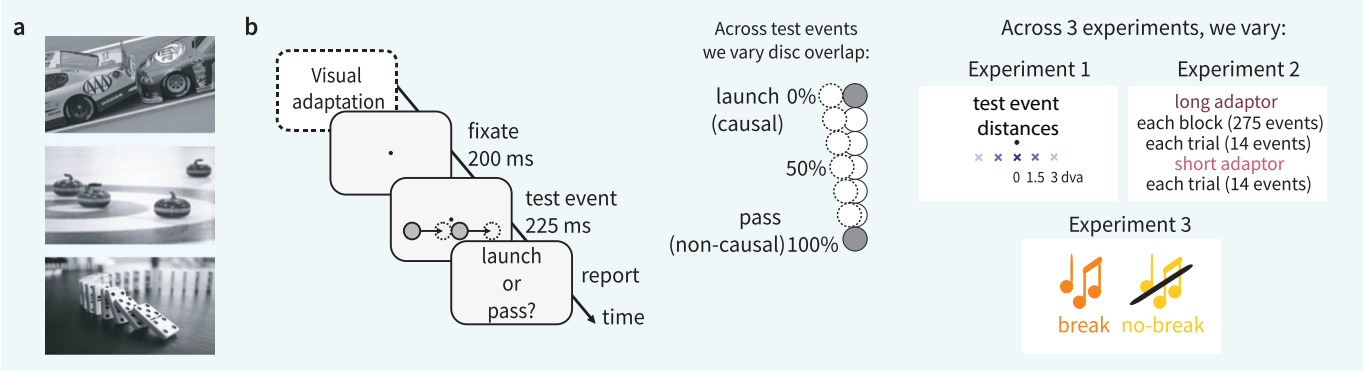


Fig. 1. Theoretical background and experimental protocol. (a) Examples of causal interactions. (b) Basic experimental paradigm, varied across three experiments. In all three experiments, we varied the disc overlap in test events from launch (causal, 0% overlap) to pass (non-causal, 100% overlap) stimuli in 7 equidistant steps. In Experiment 1, we varied the distance of the test event from the adaptor. In Experiment 2, we varied the adaptor type (long, short, or non-causal control). In Experiment 3, we introduced a break between adaptation and post-adaptation blocks.

that stops next to another disc appears to launch the second disc into motion (the *launching effect*, Fig. 1b; Michotte, 1963/2017; see Rips, 2011; Scholl & Tremoulet, 2000; White, 2006; White, 2017 for reviews). To test for spatial specificity of the perception of causal launches, we varied the center point of test events—while keeping all other characteristics constant—and determined whether visual adaptation to launches transfers to test events with a different distance from the adapted location. If the strength of the adaptation's perceptual consequences depends on the spatial location of the adaptor and the test event, it would reveal that the mechanisms for the detection of causal interactions are tightly tuned to specific spatial locations. Alternatively, if we observe adaptation transfer from one distance from the adapted location to another, this will support the notion of a broadly tuned mechanism.

To address how the perception of causality recovers from visual adaptation, we presented test events not only before and during, but also after adaptation, once we had removed the adapting stimulus. This design allowed us to assess whether the perception of causality quickly returned to baseline, which would suggest a highly flexible and rapid visual mechanism for the detection of causal interactions. In three experiments, we assessed how spatially specific visual adaptation is, whether the perception of causality completely recovers from adaptation, and whether recovery is task-based or time-based. To preview our results: we provide compelling evidence for highly spatially specific computations underlying the perception of causality, instantaneously responding to the onset of adaptation and then also quickly recovering from visual adaptation even without task-relevant input.

2. Methods

2.1. Observers and materials

2.1.1. Observers

We tested 12 human observers in Experiment 1 (ages 22–34 y; 6 female; 6 male; 12 right-handed; 6 right-eye dominant), 14 in Experiment 2 (ages 18–34 y; 7 female; 7 male; 10 right-handed; 10 right-eye dominant), and 12 in Experiment 3 (ages 18–26 y; 9 female, 3 male; 10 right-handed; 5 right-eye dominant). There were four sessions in Experiments 1 and 2 (one training session without adaptation, three test sessions with adaptation), and three in Experiment 3 (one training session, two test sessions with adaptation). The current sample size was comparable with previous work investigating the perception of causality (e.g., Ohl & Rolfs, 2025; Sommer, Rolfs, & Ohl, 2025). Based on the data obtained in the session without adaptation (i.e., the first session), we assessed whether observers can distinguish between passes and launches in the basic experiment by identifying whether the proportion of reported passes increases with increasing disc overlap. From Experiment 2, two observers were excluded from the final analyses. These two observers successfully completed the training phase and initially met the performance criterion, but in the data from the test sessions their extreme response profiles did not allow us to extract PSEs from their psychometric curves. Data from the first session (that is, without adaptation) did not enter the final analyses. We paid observers €10 per session as compensation for participation. Students enrolled in a Psychology Bachelor's programme could opt for course credits or a combination of money and credits as compensation. We obtained observers' written informed consent before the first session. All observers had normal or corrected-to-normal vision. The study was approved by the ethics committee of the Department of Psychology at the Humboldt-Universität zu Berlin, Germany, and followed the guidelines of the Declaration of Helsinki (‘World Medical Association Declaration of Helsinki’, 2013).

2.1.2. Materials

Observers sat in a sound- and brightness-attenuated room putting their head on a chinrest. 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. The eye tracker was calibrated (9 points) at the beginning of a session, after breaks, and whenever necessary throughout the experiment. 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 × 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 run on a DELL Precision T7810 (Debian GNU Linux 8) and implemented in Matlab 2023b (Mathworks, Natick, MA, USA) using the Psychophysics Toolbox 3 (Brainard, 1997; Kleiner, Brainard, Pelli, & Ingling, 2007; Pelli, 1997) for stimulus presentation and the Eyelink Toolbox (Cornelissen, Peters, & Palmer, 2002) for control of the eye tracker. Behavioral responses were collected by registering one of the two possible keypresses on a standard keyboard. In Experiment 3, an Apple iPod Nano (5th Generation; Apple Inc., released 2008) with standard wired in-ear headphones was used to present audio.

2.2. General procedure

2.2.1. Task instructions

In all 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 (Fig. 1b). 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 instructions that were displayed on the screen. As part of the instructions, we presented demo trials of test events with 0% overlap that are typically perceived as launches and test events with 100% disc overlap that are typically perceived as passes. Observers had the opportunity to inspect these two events as often as they wanted. Following the instructions, we ran a short training block of 35 (Experiment 1) or 28 trials (Experiment 2 and 3) with 7 different disc overlaps ranging from 0 to 100% overlap in equidistant steps.

2.2.2. Trial structure

At the beginning of a trial, we asked observers to fixate a gray fixation point with a diameter of 0.2 degrees of visual angle (dva) in the center of the screen on a black background. Trials started once observers successfully fixated for at least 200 ms. We presented the test events at 3 dva below the fixation point. In Experiment 1, the center of the test event varied between 0, ±1.5, and ±3 dva along the horizontal axis (relative to the adapted location). In test events, the first moving disc (gray; diameter of 1.5 dva) located either left or right from the vertical meridian started moving toward a stationary disc (gray; diameter of 1.5 dva). The first disc stopped moving at one of seven possible distances away from the stationary disc, resulting in 7 different disc overlaps ranging from 0 to 100% in equidistant steps. The entire duration of the test event was 225 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 or the arrow down key.

2.2.3. Experimental sessions

The first session served to determine whether observers perceived launches and passes as a function of disc overlap. To this end, in Experiment 1, observers completed a training session with 10 blocks in which we varied the horizontal center point of the test events (within-subjects: 5 horizontal center distances at 0, ±1.5, and ±3 dva from the adapted location) and the amount of disc overlap (within-subjects: 7 disc overlaps ranging from 0 to 100% overlap, in equidistant steps). We presented these conditions in a randomized order within a block. Each combination of these manipulations was presented once in a block, resulting in 35 trials in each block and a total of 350 trials in the first session. In Experiments 2 and 3, observers also completed 10 blocks in

the first session, but the horizontal center point of the test events was always at 0 dva from the adapted location. Instead, every combination of the disc overlap manipulation was repeated four times in each block, resulting in 28 trials in a block and a total of 280 trials in the first session. In all experiments, we counterbalanced the direction of the adaptation and test event (left-to-right vs. right-to-left) and the mapping of the response keys to reflect 'launch' or 'pass' (arrow up vs. arrow down key) between observers. In sessions 2–4 of Experiment 2, we presented three different adaptor types that varied between the sessions. In sessions 2 and 3 of Experiment 3, we introduced a break between adaptation and post-adaptation blocks in either of the two sessions. See details of the procedures for the specific experiments below.

2.2.4. Eye-tracking

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 the eyes' deviation from the fixation point online. We aborted a trial whenever this distance 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) saying "please fixate" if their gaze exceeded 2 dva away from the fixation point.

2.3. Adaptation in Experiment 1

In Experiment 1, each observer completed 18 blocks in each of sessions 2–4. The first six blocks in a session were without adaptation to measure an observer's perception of causality before adaptation. As in the first session, we presented test events at five possible horizontal center locations and varied the disc overlap, resulting in 35 trials in each block. In blocks 7–12, we presented an adaptor before the first trial of a block. The adaptor was a stream of launching events in the same direction as the test event. To this end, we presented 275 launching events before the first trial of a block, each located at the same horizontal offset as the fixation point (0 dva) and vertically at 3 dva below the fixation point. The exact directions of these launching events were randomly chosen from a narrow uniform distribution around the direction on the horizontal meridian (±30 degrees) to avoid adaptation to low-level features (such as contrast and luminance) rather than specifically to the causal percept (Ohl & Rolfs, 2025). The long adaptation phase was complemented by a top-up adaptation of 14 launching events in the same direction as the adaptor (drawn from the same narrow uniform distribution around the main direction) before each trial to keep adaptation effective throughout the entire block. Blocks 13–18 were without adaptation, so that we could address the recovery from visual adaptation. Each observer completed a total of 1890 trials in Experiment 1.

2.4. Variations in Experiment 2

In Experiment 2, we presented three different adaptor types to address the strength of the recovery of visual adaptation. Each of sessions 2–4 consisted of a different type of adaptor: one which was exactly as in Experiment 1, which combined the long adaptation sequence of 275 launching events at the start of each adaptation block with 14 top-up events in each adaptation trial ('long' adaptation), one which only had the 14 top-up events ('short' adaptation) and a non-causal control. The purpose of the short adaptor was to see whether it would elicit the same adaptation response as a longer adaptor. For the control condition, we used non-causal 'slip' events (Rolfs et al., 2013) as adaptors to control that the adaptation was specific for the impression of causality in the launching events. These control events matched launching events in as many physical properties as possible while producing a very different, non-causal phenomenology. In non-causal control events, the first disc also moved toward a stationary second disc. In contrast to launching events, however, the first disc passed the stationary disc and stopped only when it was 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 adaptors, we displayed control 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 per observer). This experimental design resulted in three test sessions of which the order was randomly determined for each observer. This time, blocks 13–22 were without adaptation, so that we could address the recovery from visual adaptation over a longer time interval as compared to Experiment 1 (6 blocks of post-adaptation in Experiment 1 vs. 10 post-adaptation blocks in Experiment 2). Each observer completed a total of 1848 trials in Experiment 2.

2.5. Variations in Experiment 3

In Experiment 3, we presented only top-up adaptors, as our results from Experiment 2 suggested that short initial adaptation events sufficed to elicit the same magnitude of adaptation as a long initial adaptation event. In Experiment 3, we aimed to address whether we recover from visual adaptation in a task-relevant manner by actively recalibrating the visual system to the statistical regularities in the environment, or whether this is a more passive, time-based process. That is, do we recover from adaptation by seeing numerous non-causal events in the environment? To test this, observers completed two sessions that differed only in when (following block 12) the adaptation phase ended. In the no-break condition, observers proceeded directly from the adaptation blocks to the post-adaptation blocks, replicating the continuous timeline used in Experiment 2. In the break condition, a 10-min interval without task-relevant visual input was inserted before the post-adaptation blocks, during which observers listened to a podcast. During this break, observers were required to stay seated and listen to one of three podcast options: "The Science of Birth Without Pregnancy" (O'Neill, 2021), "The Science of Designer Babies" (O'Neill, 2024, November 7), or "The Science of Heart Attacks" (O'Neill, 2024, November 28). The purpose of the podcast was to present entertainment whilst simultaneously not presenting any visual causal interactions of any kind. This allowed us to address whether the perception of causality also recovered when there was no task-relevant input that would suggest different statistical regularities of the environment, while time did pass. The podcast break was exactly 10-min long, matching the average time the post-adaptation blocks took in Experiment 2. Thus, the break condition was directly comparable to the first approximately 10-min of the post-adaptation phase in Experiment 2, except that observers received no task-relevant visual stimulation during this interval. The observer's eyes were not tracked during the podcast break. This experimental design resulted in two test sessions, one with and one without a break, of which we counterbalanced the order for each observer. Like in Experiment 2, blocks 13–22 were without adaptation, so that we could address the recovery from visual adaptation over a longer time interval compared to Experiment 1 (i.e., 10 rather than 6 post-adaptation blocks). Each observer completed a total of 1232 trials in Experiment 3.

2.6. Data analysis

For the statistical analyses and estimation of the psychometric functions, we used the *quickpsy* package (Linares & Lopez-Moliner, 2016) and the RStudio environment (Posit team, 2025). 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 of 0–0.75, and the upper asymptote in the range of 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.

To investigate patterns of the perception of causality block-by-block and over the different phases of the experimental sessions (before, during, and after adaptation), we related causal reports over blocks and phases to the three different testing distances from the adapted location (0 dva = center; ±1.5 dva = near; ±3 dva = far) in Experiment 1, the three different adaptors (long, short, control) in Experiment 2, and the two break conditions (podcast-break or no-break) in Experiment 3. Again, we calculated the mean proportions separately for each observer and condition and then calculated group means. To confirm that symmetrical distances could be collapsed in Experiment 1, we statistically compared the psychometric thresholds at each distance (−1.5 vs +1.5 dva; −3 vs +3 dva) using paired *t*-tests. No significant differences were found for either the near or far conditions in any phase (all $ps \geq 0.19$),

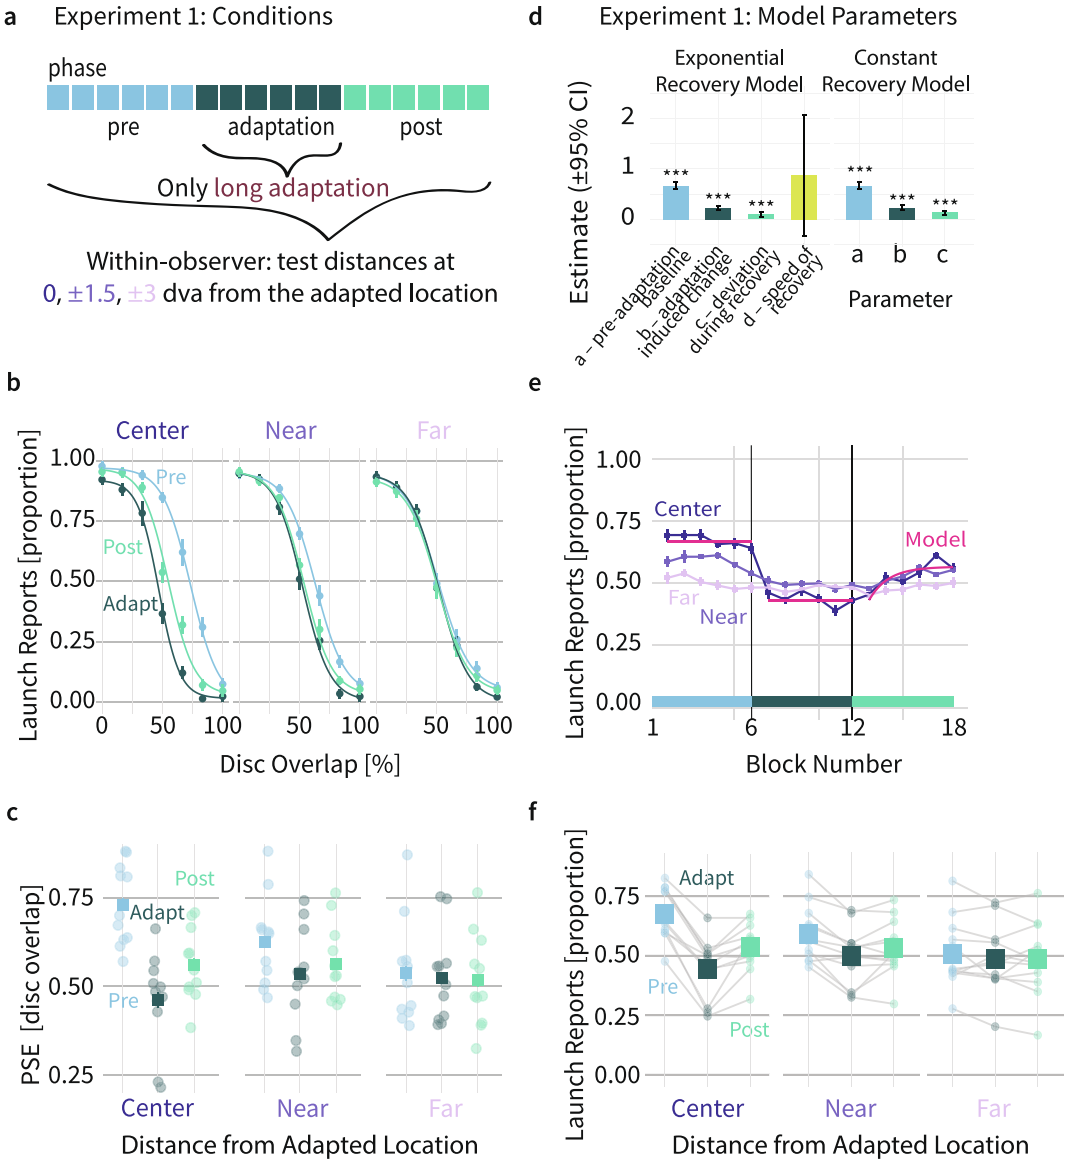


Fig. 2. Results of Experiment 1. (a) Experiment 1 consisted of 18 blocks, divided into pre-adaptation, adaptation, and post-adaptation phases. The adaptor consisted of 275 events at the start of each block, and 14 events at the start of each trial. Test events were presented at horizontal distances at 0 (center), ±1.5 (near), and ±3 dva (far) from the adapted location. (b) Mean proportion of launch reports as a function of disc overlap. Visualization of psychometric curves is based on the mean parameters averaged across observers before adaptation (light blue), during adaptation (dark green), and after adaptation (mid green) for the different test event locations: center (dark purple), near (mid purple) and far (light purple). (c) PSEs for individual observers (circles, $n = 11$) and the mean across observers (square) for the different test event locations. (d) Estimated parameters for competing exponential and constant piecewise non-linear mixed effects models. Model fitting was restricted to trials at the center location (0 dva distance), as this was the only location showing reliable adaptation and recovery effects. Parameter a denotes the pre-adaptation baseline, b denotes the adaptation-induced change relative to the baseline, c denotes the asymptotic deviation from baseline during recovery, and d denotes the speed of the recovery. For each of the parameters, we denote 95% confidence intervals. $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$. (e) Blockwise analysis of mean proportion of launch reports. Visualization of blockwise averages is based on the proportion of launch reports per block in the experiment, averaged across observers for each test event location. Additionally, the winning non-linear mixed effects model fit (exponential model, bright pink) for trials at the center distance (0 dva). (f) Launch reports for individual observers (circles, $n = 12$) and the mean across observers (square) for the different phases. Error bars indicate Cousineau–Morey 95% confidence intervals (Morey, 2008). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

indicating that symmetric positions produced equivalent PSEs. Therefore, we collapsed ±1.5 dva into a single "near" level and ±3 dva into a "far" level.

For inferential statistics, we analyzed PSEs using repeated-measures analyses of variance (rmANOVA). Observers were included in the PSE analysis only if valid threshold estimates were obtained in all conditions. Based on this criterion, one observer was excluded in Experiment 1 and two observers were excluded in Experiment 2. Error bars indicate Cousineau–Morey 95% confidence intervals (Morey, 2008). A significant interaction in the rmANOVA was followed up with planned post-hoc pairwise comparisons of estimated marginal means using uncorrected paired contrasts.

The model comparisons for non-linear mixed effects models were based on Akaike and Bayesian Information Criteria (AIC, BIC) computed from maximum-likelihood fits.

Experiments 1 and 3 were preregistered on the Open Science Framework (OSF), and data and analysis code for all experiments are publicly available (https://osf.io/e4yjn).

3. Results

In three experiments, we tested how the perception of causality recovers from visual adaptation. Previous adaptation studies on causality perception have typically compared the proportion of launches (or points of subjective equality, PSEs) before adaptation to those measured during adaptation. To our knowledge, no study has additionally compared the proportion of launches after adaptation to the pre-adaptation and adaptation phases. Here, we evaluated recovery by directly contrasting the PSEs after adaptation to those measured before and during adaptation.

3.1. Adaptation to causal events is spatially specific

In Experiment 1, we examined the points of subjective equality (PSEs) across the three phases (pre, adaptation, post) as a function of test event locations (center, near, far) using a two-way repeated-measures ANOVA, followed by pairwise contrasts. Moreover, we fitted non-linear mixed effects models designed to answer whether there was an adaptation effect (pre-adaptation vs. adaptation), if so, whether recovery occurred, (post-adaptation vs. adaptation), and if so, whether the recovery was complete (post-adaptation vs. pre-adaptation). We additionally tested whether adaptation to launch events is spatially specific to the center of the adapting stimulus, or whether it spreads to nearby test event locations. To this end, we presented test events at one of three possible horizontal distances from the adapted location (0, ±1.5, and ±3 dva; Fig. 2a). The first disc of the test event stopped as soon as the two partially overlapped (manipulated in 7 equidistant steps from no to full overlap). The second disc then started to move in the same direction as the first disc, and observers reported whether they perceived that the first disc launched the second (common at no overlap) or passed it (common at full overlap). We quantified the PSE between perceived launches and passes by modeling psychometric functions that related disc overlap to the proportion of launch reports before, during, and after adaptation.

Visual adaptation reduced the perception of causality at the center and near test event locations, but not for the far location (Fig. 2b,c). Observers reported fewer launches during adaptation compared with pre-adaptation, with the strongest effect at the center ($PSE_{pre} = 0.729 \pm 0.038$; $PSE_{adaptation} = 0.462 \pm 0.038$; $\Delta PSE = 0.267 \pm 0.025$), weaker effects at the near location ($PSE_{pre} = 0.625 \pm 0.038$; $PSE_{adaptation} = 0.534 \pm 0.038$; $\Delta PSE = 0.091 \pm 0.025$), and no adaptation at the far location ($PSE_{pre} = 0.539 \pm 0.038$; $PSE_{adaptation} = 0.524 \pm 0.038$; $\Delta PSE = 0.015 \pm 0.025$). After adaptation, PSEs partially recovered during the post-adaptation phase (Fig. 2c; center: $PSE_{post} = 0.562 \pm 0.038$; $\Delta PSE = 0.099 \pm 0.025$; near: $PSE_{post} = 0.563 \pm 0.038$; $\Delta PSE = 0.028 \pm 0.025$; far: $PSE_{post} = 0.517 \pm 0.038$; $\Delta PSE = -0.007 \pm 0.025$).

A two-way rmANOVA (3 phases: pre-adaptation, adaptation, post-adaptation; 3 distances from the adapted location: center, near, far location) on the PSEs revealed a significant main effect of phase ($F(2,20) = 19.74$, $p < .001$, $\eta p^2 = 0.66$), a marginally significant effect of distance ($F(2, 20) = 3.28$, $p = .059$, $\eta p^2 = 0.25$), as well as a significant interaction between phase and distance ($F(4, 40) = 27.70$, $p < .001$, $\eta p^2 = 0.73$).

A one-way rmANOVA on pre-adaptation PSEs revealed a significant effect of distance ($F(2, 20) = 17.92$, $p < .001$, $\eta p^2 = 0.64$). Pairwise comparisons showed that PSEs were highest at the center location, intermediate at the near location, and lowest at the far location. Specifically, PSEs at the center were higher than at the near ($\Delta = 0.104$, $p < .001$) and far locations ($\Delta = 0.190$, $p < .001$), and PSEs at the near location were higher than at the far location ($\Delta = 0.086$, $p < .001$), indicating that baseline causal perception varies systematically with test event distance from the adapted location.

Estimated marginal mean contrasts confirmed the spatial specificity of the adaptation effect. First, pre vs. adaptation (testing the adaptation effect) showed a large reduction at the center (estimate $= 0.267 \pm 0.025$, $p < .001$), a smaller but reliable effect at the near location (estimate $= 0.091 \pm 0.025$, $p < .001$), and a very small effect at the far location (estimate $= 0.015 \pm 0.025$, $p = .551$). Pairwise contrasts further showed that the adaptation effect at the center was larger than at near and far, and near was larger than far ($p < .05$). Second, post- vs. adaptation comparisons (testing recovery) revealed significant recovery at the center (estimate $= 0.099 \pm 0.025$, $p < .001$), but not at the near (estimate $= 0.028 \pm 0.025$, $p = .247$) or far locations (estimate $= -0.007 \pm 0.025$, $p = .571$).

3.2. Recovery from adaptation is rapid but gradual

To characterize the temporal dynamics of the recovery from adaptation in more detail, we modeled responses across blocks using a piecewise non-linear mixed-effects model. This model consisted of three pieces: a constant pre-adaptation baseline, a constant adaptation level, and an exponential recovery function during the post-adaptation phase. Responses $y(x)$ as a function of block number x were modeled as

$$y(x) = \begin{cases} a, x \le 6 \text{ (pre} - \text{adaptation)} \\ a - b, 7 \le x \le 12 \text{ (adaptation)} \\ (a-b) + (b-c)\left(1 - e^{-d(x-13)}\right) x \ge 13 \text{ (post} - \text{adaptation)} \end{cases}$$

where a denotes the pre-adaptation baseline, b denotes the adaptation-induced change relative to the baseline, c denotes the asymptotic deviation from baseline during recovery, and d denotes the speed of the recovery. This model, assuming a gradual recovery via the exponential function, we then compared to a reduced model assuming an instantaneous, step-like recovery. This reduced model omitted the exponential term and assumed a constant post-adaptation level $a - c$. Random effects were included on parameters a, b, and c to account for between-observer variability. Both models were restricted to trials at the center location, as this was the only location showing reliable adaptation and recovery effects; near and far locations were therefore excluded from this analysis.

Model comparison favored the exponential model, which showed lower information criteria (AIC = 5234.26, BIC = 5302.97) than the reduced constant recovery model (AIC = 5241.44, BIC = 5303.90). While BIC showed only a marginal preference, AIC provided support for gradual recovery dynamics. This indicates that recovery unfolded progressively over multiple post-adaptation blocks rather than occurring

abruptly at the transition from adaptation to post-adaptation (Fig. 2d–f). Together, these results are more consistent with a gradual than an instantaneous recovery profile in Experiment 1, although this conclusion is based primarily on the lower AIC of the exponential model. Importantly, despite this gradual recovery, post-adaptation did not fully return to pre-adaptation levels. Parameter c, the asymptotic deviation from baseline during recovery, was significantly greater than zero ($c = 0.099$, $SE = 0.029$, $p < .001$). Thus, the observed recovery from adaptation was incomplete in this experiment.

3.3. Recovery from adaptation with an extended post-adaptation period

In Experiment 2, we extended the post-adaptation phase compared to Experiment 1 to assess whether the perception of causality fully recovers with more time. Observers experienced three adaptor types: long, short, and a non-causal control (Fig. 3a). Points of subjective equality (PSEs) were quantified before, during, and after adaptation for the three adaptor types. As before, we tested whether adaptation occurred (pre vs. adaptation), whether recovery occurred (post vs. adaptation), and

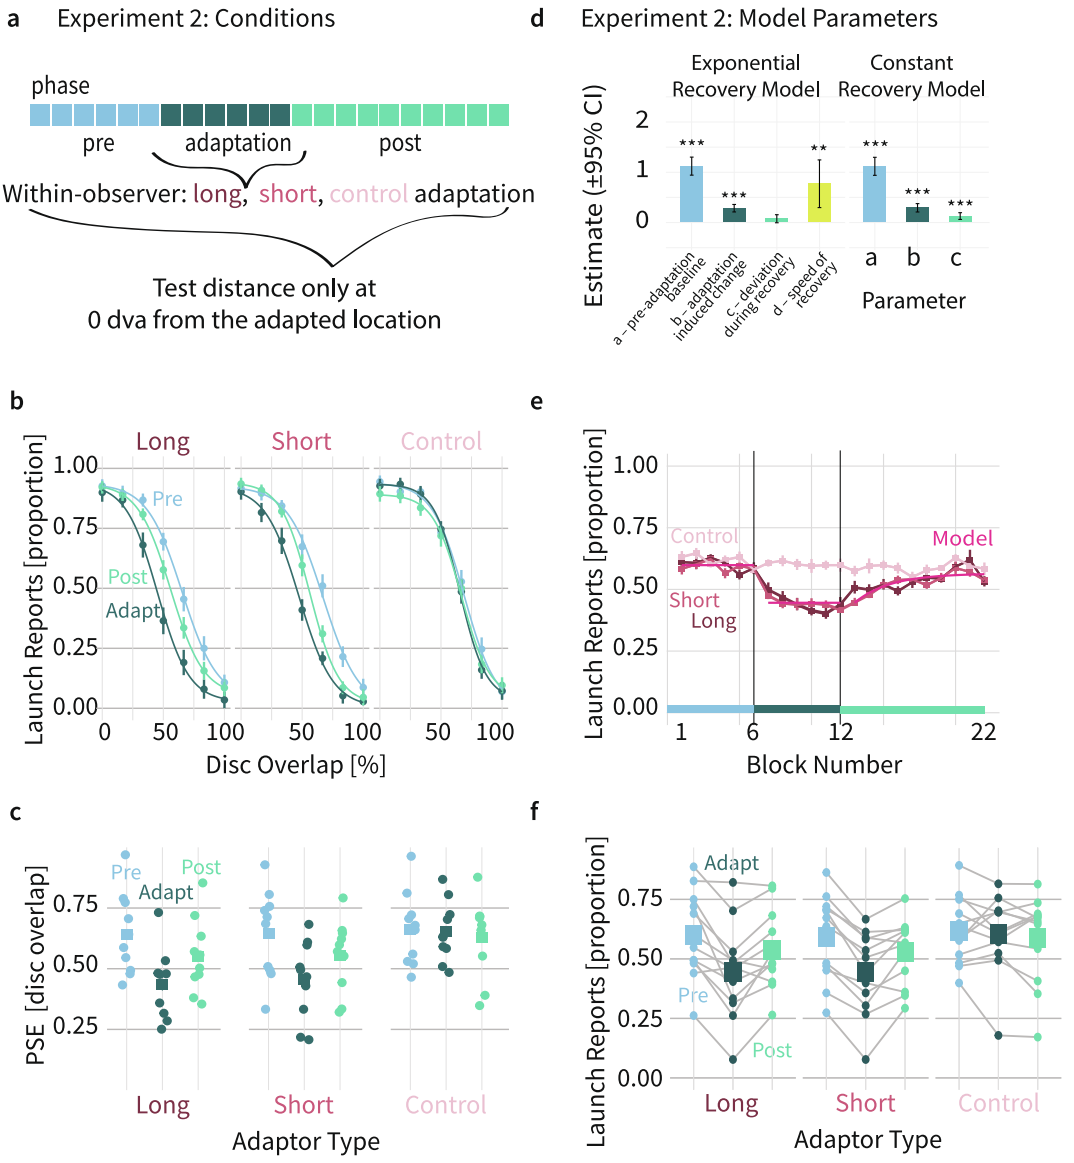


Fig. 3. Results of experiment 2. (a) Experiment 2 consisted of 22 blocks, divided into pre-adaptation (6), adaptation (6), and post-adaptation (10) phases. The long as well as the control adaptor consisted of 275 events at the start of each block, and 14 events at the start of each trial. The short adaptor consisted only of 14 events at the start of each trial. (b) Mean proportion of launch reports as a function of disc overlap. Visualization of psychometric curves is based on the mean parameters averaged across observers before adaptation (light blue), during adaptation (dark green), and after adaptation (mid green) for the different adaptor types: long (dark purple), short (mid pink), and the non-causal control (light pink). (c) PSEs for individual observers (circles, $n = 10$) and the mean across observers (square) for the different adaptor types (same order as in the above panel). (d) Estimated parameters for competing exponential and constant piecewise non-linear mixed effects models. For model fitting, trials from long and short adaptors were collapsed, since we observed no difference between the two adaptors. Parameter a denotes the pre-adaptation baseline, b denotes the adaptation-induced change relative to the baseline, c denotes the asymptotic deviation from baseline during recovery, and d denotes the speed of the recovery. For each of the parameters, we denote 95% confidence intervals. $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$. (e) Blockwise analysis of mean proportion of launch reports. Visualization of blockwise averages is based on the mean launch reports per block in the experiment, averaged across observers for the long adaptor (dark purple), short adaptor (mid pink) and non-causal control adaptor (light pink). Additionally, the winning non-linear mixed effects model fit (exponential model, bright pink) for trials from long and short adaptors. (f) Launch reports for individual observers (circles, $n = 12$) and the mean across observers (square) for the different phases. In all subfigures, error bars indicate Cousineau–Morey 95% confidence intervals (Morey, 2008). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

whether recovery was complete (post vs. pre).

During the adaptation phase, long and short adaptors produced clear decreases in causal perception (Fig. 3b,c; long: $\Delta PSE = -0.207 \pm 0.033$; short: $\Delta PSE = -0.190 \pm 0.033$), whereas the non-causal control adaptor produced no meaningful change ($\Delta PSE = -0.009 \pm 0.033$). In the post-adaptation phase, causal perception partially recovered for long and short adaptors (Fig. 3c; long $\Delta PSE = 0.117 \pm 0.033$; short $\Delta PSE = 0.099 \pm 0.033$), while the control remained near baseline ($\Delta PSE = -0.024 \pm 0.033$).

A two-way repeated-measures ANOVA with factors phase (pre, adaptation, post) and adaptor type (long, short, control) revealed significant main effects of phase ($F(2, 18) = 13.16$, $p < .001$, $\eta p^2 = 0.59$) and adaptor type ($F(2, 18) = 13.56$, $p < .001$, $\eta p^2 = 0.60$), as well as a significant interaction ($F(4, 36) = 11.02$, $p < .001$, $\eta p^2 = 0.55$), indicating that the magnitude of adaptation depended on the adaptor type.

Planned contrasts testing the adaptation effect (pre vs. adaptation) confirmed significant reductions in PSEs for both the long (estimate $= 0.207 \pm 0.033$, $p < .001$) and short (estimate $= 0.190 \pm 0.033$, $p < .001$) adaptors relative to baseline, whereas no adaptation effect was observed for the control adaptor (estimate $= 0.009 \pm 0.033$, $p = .791$). Direct comparisons showed that adaptation effects were significantly larger for the long and short adaptors than for the control (long vs. control: $t(37.6) = 6.30$, $p < .001$; short vs. control: $t(37.6) = 5.77$, $p < .001$), while the long and short adaptors did not differ from each other ($t(37.6) = 0.27$, $p = .791$). To provide additional evidence that the long and short adaptors were indeed not different, we computed a Bayesian paired *t*-test. This provided weak-to-moderate evidence for the absence of a difference ($BF01 = 2.56$).

We next assessed the dynamics of the recovery from adaptation by fitting the same two competing piecewise nonlinear mixed-effects models as in Experiment 1: a model with an exponential recovery component and a reduced model assuming instantaneous (constant) recovery. Both models were fit after excluding the slip adaptor condition; only the long and short adaptor conditions were included in the modeling. Model comparison revealed a modest advantage for the exponential model, which provided a lower AIC than the constant model (19,914.7 vs. 19,917.2) (Fig. 3d–f). However, BIC values were very similar between the two models (19,998.4 vs. 19,993.3). Therefore, we caution for an overinterpretation of the evidence for the exponential model, noting that while a gradual recovery is plausible, the data do not rule out a simpler constant-recovery account. Finally, to assess whether recovery was complete, we compared post-adaptation reports to pre-adaptation baseline by investigating parameter *c*, the asymptotic deviation from baseline during recovery. The parameter was not significantly different from zero ($c = 0.033$, $SE = 0.022$, $p = .135$), suggesting that post-adaptation reports returned to pre-adaptation baseline. Recovery after adaptation was therefore complete in Experiment 2.

Together, the results of Experiment 2 show that adaptation to causal launch events robustly reduces perceived causality, that this reduction weakens once the adaptor is removed, and that extending the post-adaptation phase is most compatible with a gradual and complete recovery profile, although the data did not clearly distinguish this from a more parsimonious constant-recovery account. Importantly, recovery dynamics did not differ between long and short adaptors, and no adaptation or recovery was observed for a non-causal control adaptor, mirroring the pattern of specificity observed in Experiment 1.

3.4. Recovery from adaptation after a 10-min break

In Experiment 3, we asked whether the perception of causality recovers in the absence of experimentally controlled task-relevant visual input, that is, purely as a function of time. To this end, observers either took a 10-min "podcast-break" between the adaptation and post-adaptation phase (viewing a neutral display), or had no break between the two phases (Fig. 4a). We again quantified PSEs before, during, and after adaptation to assess if there was an adaptation effect (pre vs. adaptation), and if so, whether recovery occurred (post vs. adaptation).

PSEs were nearly identical across the two break conditions in all three phases (Fig. 4b,c; pre: $PSE_{break} = 0.631 \pm 0.045$; $PSE_{no\text{-}break} = 0.619 \pm 0.045$; adaptation: $PSE_{break} = 0.466 \pm 0.045$; $PSE_{no\text{-}break} = 0.442 \pm 0.045$; post: $PSE_{break} = 0.571 \pm 0.045$; $PSE_{no\text{-}break} = 0.577 \pm 0.045$). A two-way repeated-measures ANOVA (3 phases: pre, adaptation, post, 2 break types: podcast-break, no-break) revealed a significant main effect of phase ($F(2,22) = 20.85$, $p < .001$, $\eta p^2 = 0.65$), indicating that adaptation reliably reduced PSEs compared to the pre-adaptation phase. Neither the main effect of break type ($F(1,11) = 0.61$, $p = .451$, $\eta p^2 = 0.05$) nor the interaction between phase and break type ($F(2, 22) = 1.26$, $p = .302$, $\eta p^2 = 0.10$) was significant, indicating that a 10-min podcast break did not reliably modulate adaptation or recovery. Thus, both the podcast-break and no-break conditions showed fast and comparable recovery, with no evidence that inserting a 10-min break delayed the recovery.

As in Experiment 1 and Experiment 2, we modeled the temporal profile of the recovery from adaptation by fitting the same two competing piecewise nonlinear mixed-effects models. Since both break types showed similar adaptation and recovery effects, we collapsed the two conditions for model fitting. As before, we compared a model with an exponential recovery component to a reduced model assuming instantaneous (constant) recovery. In contrast to Experiment 1 and Experiment 2, the exponential model failed to converge for Experiment 3, despite multiple starting values and reduced random-effects structures. Because model non-convergence can arise for several reasons, including parameter identifiability and limited information in the data, we do not interpret this result as evidence against gradual recovery. Nevertheless, the constant model converged reliably ($AIC = 20{,}366.4$, $BIC = 20{,}442.4$), and provided an adequate description of the observed data (Fig. 4d–f). Thus, recovery in Experiment 3 was best characterized as instantaneous rather than gradual. Finally, to assess whether recovery was complete, we compared post-adaptation reports to pre-adaptation baseline by investigating parameter *c*, the asymptotic deviation from baseline during recovery. The parameter remained significantly above zero ($c = 0.039$, $SE = 0.013$, $p < .01$), suggesting that post-adaptation reports did not fully return to baseline. Thus, recovery in Experiment 3 was instantaneous but incomplete, consistent with a rapid offset of the adaptation effect followed by a persistent residual bias.

Together, these results show that in Experiment 3, the negative aftereffect in perceived causality weakens over time, but recovery remains incomplete even after a 10-min interruption. Importantly, the temporal break did not facilitate recovery relative to continuous viewing, indicating that decay over time alone was insufficient to fully reverse adaptation to causality.

4. Discussion

We set out to investigate the spatial specificity and temporal dynamics of visual adaptation to causal events. Across three experiments, we found that the effects of visual adaptation were narrowly tied to the spatial location of the test event. Recovery occurred rapidly in all three experiments, although the precise recovery profile differed across experiments. Experiments 1 and 2 were more consistent with gradual recovery, whereas the data from Experiment 3 were more adequately captured by a parsimonious constant-recovery model. In all three experiments, causality perception returned toward baseline levels within the same testing session. We replicated the general adaptation effect reported in previous studies (e.g., Kominsky & Scholl, 2020; Kominsky & Wenig, 2025; Ohl & Rolfs, 2025; Rolfs et al., 2013), showing that visual adaptation reliably alters the perception of causality. Although the present experiments do not distinguish between the mechanisms underlying this adaptation effect, they characterize key spatial and temporal properties that place constraints on possible accounts. In Experiment 1, we observed a negative aftereffect during adaptation that varied with distance from the adapted location: adaptation effects were

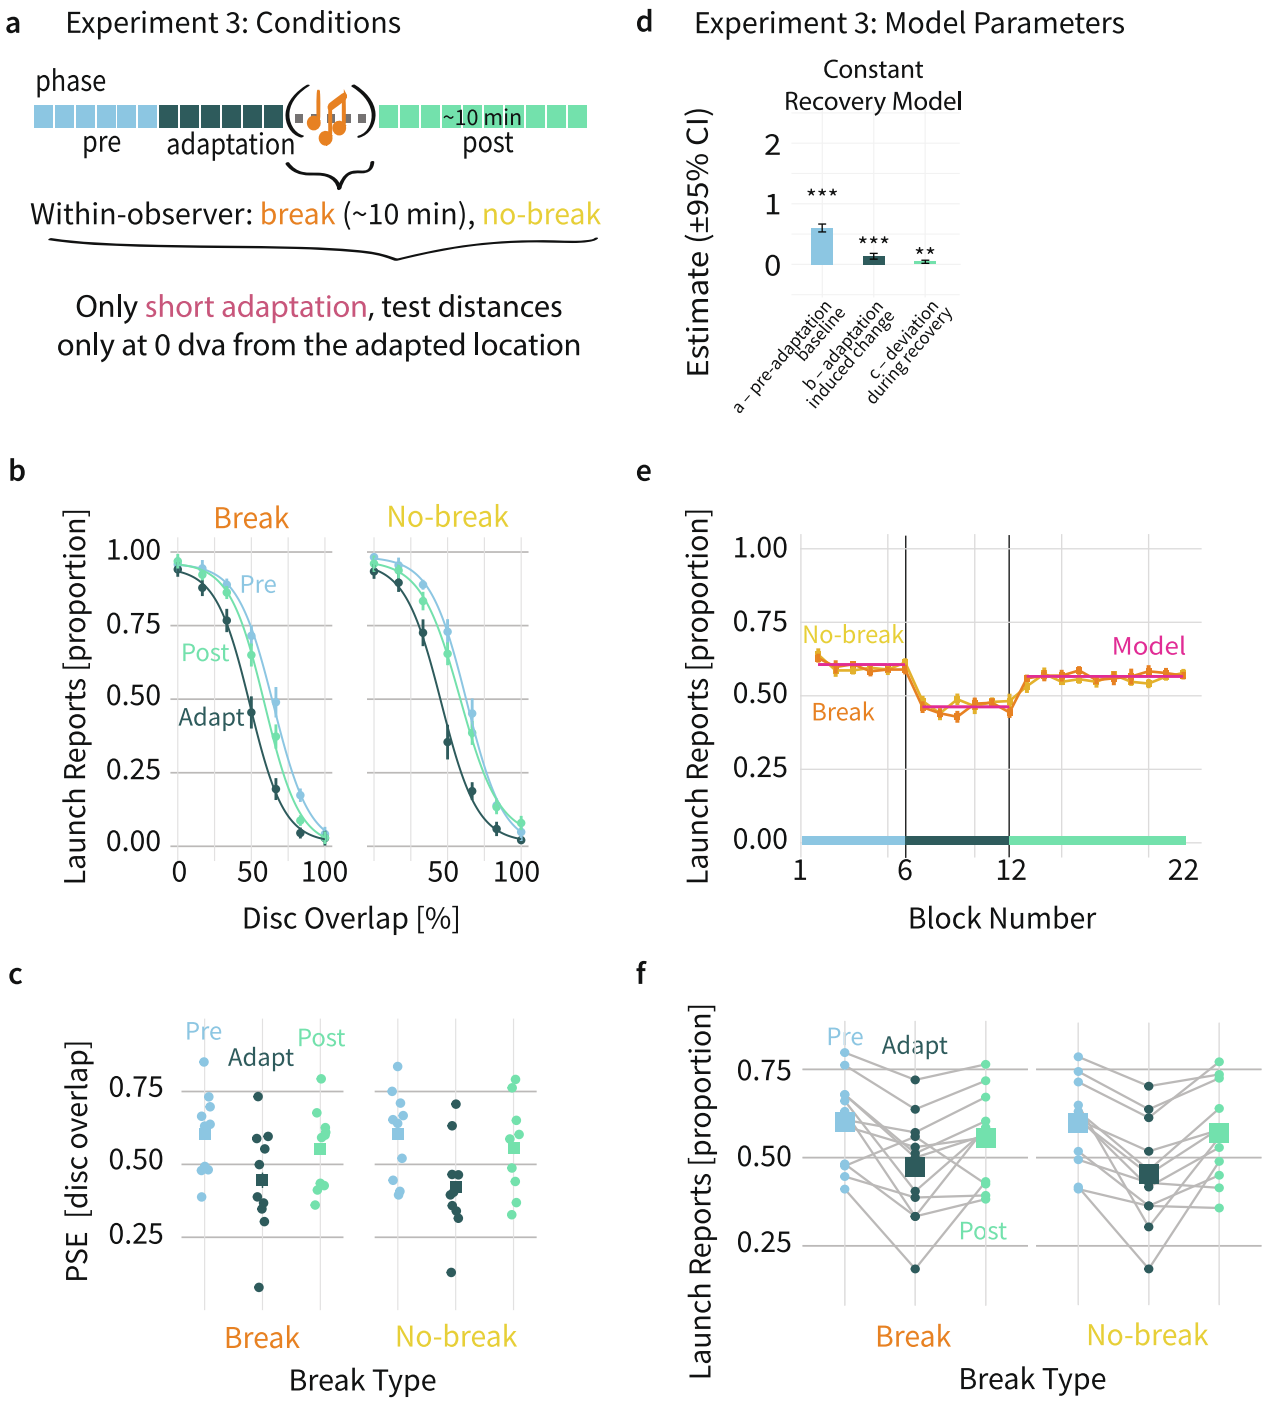


Fig. 4. Results of Experiment 3. (a) Experiment 3 consisted of 22 blocks, divided into pre-adaptation, adaptation, and post-adaptation phases. The adaptor consisted of 14 events at the start of each trial. Observers took part in two testing sessions, one with a podcast break between the adaptation and post-adaptation phases, and one session without a break. (b) Mean proportion of launch reports as a function of disc overlap. Visualization of psychometric curves is based on the mean parameters averaged across observers before adaptation (light blue), during adaptation (dark green), and after adaptation (mid green) for the different break conditions: break (orange), and no-break (yellow). (c) PSEs for individual observers (circles, $n = 12$) and the mean across observers (square) for the different break conditions (same order as in the above panel). (d) Estimated parameters for the constant piecewise non-linear mixed effects models. Parameter a denotes the pre-adaptation baseline, b denotes the adaptation-induced change relative to the baseline and c denotes the asymptotic deviation from baseline during recovery. For each of the parameters, we denote 95% confidence intervals. $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$. (e) Blockwise analysis of mean proportion of launch reports. Visualization of blockwise averages is based on the mean launch reports per block in the experiment, averaged across observers for the break (orange) and no-break conditions (yellow). Additionally, the winning non-linear mixed effects model fit (constant model, bright pink). (f) Launch reports for individual observers (circles, $n = 12$) and the mean across observers (square) for the different phases. In all subfigures, error bars indicate Cousineau–Morey 95% confidence intervals (Morey, 2008). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

strongest at the adapted location and diminished with increasing distance. Moreover, the perception of causality gradually recovered once we removed the adaptor, but was incomplete within the tested timeframe. In Experiment 2, we extended this finding by showing that the recovery was complete with slightly more time. Additionally, a shorter adaptation duration was equally effective in driving the adaptation effect. In contrast to the gradual recovery observed in Experiments 1 and 2, Experiment 3 showed instantaneous recovery that progressed even without task-relevant input.

4.1. Spatial specificity of visual adaptation

Since adaptation was tightly confined to the spatial location of the test event, the processes contributing to the observed adaptation effects are likely to depend on mechanisms with relatively small receptive fields. Small receptive fields are a hallmark of early visual areas, where neural tuning is tightly linked to precise spatial locations (Wandell et al., 2007). By contrast, later visual areas typically integrate information across larger regions of space, supporting position-invariant representations. This is consistent with the idea that even seemingly high-level constructs like causality may be grounded in visual perceptual mechanisms (Michotte, 1963/2017; Scholl & Tremoulet, 2000). Thus, the observed spatial specificity is broadly consistent with an involvement of relatively early visual processing. However, the present behavioral data do not allow us to identify the neural locus of the adaptation effect, and any assignment to specific visual areas remains necessarily speculative. Rather, our findings provide constraints on candidate mechanisms that future physiological or neuroimaging studies will need to test.

An alternative interpretation is that the observed aftereffects reflect adaptation to lower-level spatiotemporal event structure rather than to causal representations themselves. However, Experiment 2 argues against a purely spatiotemporal account. The slip adaptor preserved many of the critical spatiotemporal characteristics of the launch event, including motion energy, temporal contiguity, and motion onset and offset, while largely eliminating the percept of causality. Unlike the launch adaptor, this condition produced little evidence for adaptation. This suggests that adaptation is not driven by spatiotemporal event structure alone, but instead depends on events that are interpreted as launch-like interactions. Importantly, this interpretation does not imply adaptation to causality in a strictly high-level sense. Rather, the observed aftereffects may reflect adaptation to structured event representations that are constrained by both low-level spatiotemporal information and the perceptual interpretation that an event is launch-like.

Our finding of high spatial specificity has important implications for how visual adaptation to causality should be interpreted. If the effect were driven by cognitive factors such as a shift in a decision criterion, that is, a change in the internal threshold for judging whether an event is a collision, then this criterion should, in principle, apply globally across the visual field. A collision is a collision, regardless of where it occurs. However, the adaptation effect we observed was *highly* spatially specific, which is difficult to reconcile with a purely criterion-based account. Although we cannot fully rule out the possibility of a spatially specific decision criterion, such a mechanism would be a rather unlikely explanation for the results in the present study. Instead, the strong spatial specificity is more consistent with the idea that our experimental manipulation altered a visual, perceptual representation itself rather than a purely global decision threshold. Recent work has highlighted that spatial specificity alone is not sufficient to establish the perceptual nature of an adaptation effect, as spatially specific adaptation can also arise for non-visual attributes, such as value, when represented in non-retinotopic reference frames (Clarke & Yousif, 2026). The retinotopic specificity reported earlier (Rolfs et al., 2013), and the high degree of spatial specificity observed in the current experiments therefore provide a more constrained link to visual perceptual mechanisms. However, the present data cannot fully exclude alternative explanations, including spatially specific decisional recalibration, attentional mechanisms, or predictive processes that are themselves organized retinotopically. Thus, while the observed spatial specificity places important constraints on candidate mechanisms, it does not uniquely identify the level of processing at which adaptation occurs.

One potential apparent discrepancy with previous work concerns the absence of nonspecific or spatiotopic adaptation effects in our data. In particular, Kominsky and Scholl (2020) and Kominsky and Wenig (2025) reported reliable nonspecific effects alongside stronger retinotopically specific effects, even when test locations were substantially displaced in space. Our finding that adaptation effects vanished with relatively modest increases in distance from the adapted location may therefore appear inconsistent with this earlier literature. However, an important distinction is that previous studies examined spatial specificity across locations far apart from each other (e.g., opposite visual hemifields), whereas the present study probed the spatial scale of adaptation over small spatial separations. By including test locations at distances of $\pm 1.5^\circ$ and $\pm 3^\circ$ relative to an adaptor on the vertical meridian, our design allowed us to examine the spatial extent of adaptation at finer spatial scales than those examined in previous work. Thus, whereas previous findings established that causal adaptation is spatially specific across hemifields, the present findings further characterize how rapidly adaptation declines with increasing distance from the adapted location. A key difference lies in the interaction between a test event's spatial location and baseline perceptual reports at that location. As eccentricity from fixation increases (smallest for center, largest for far locations), observers' pre-adaptation reports are systematically biased toward more pass reports, reflecting a shift in the underlying psychometric function rather than a loss of discriminability. Crucially, although this baseline bias toward pass percepts reduces the dynamic range available for detecting further adaptation-induced shifts in the same direction, it does not eliminate it. The psychometric functions show that observers continued to report both launch and pass percepts across the full range of overlaps, including at the largest distance from the adaptor, which also had the largest eccentricity from fixation. Thus, negative aftereffects remained measurable in principle, even if the eccentricity-dependent baseline shift reduced the maximum observable magnitude of the adaptation effect.

4.2. Recovery from visual adaptation

An important contribution of the present work is that it highlights the value of examining not only the onset of visual adaptation, but also its recovery. This recovery period offers a largely untapped time window in which to learn about the dynamics of the visual system. In our case, we observed recovery on a relatively fast timescale compared to the more commonly studied long-term examples, yet within our study this recovery was still more gradual than one might expect from such a brief adaptation phase. Notably, we found that the same adaptation and recovery patterns emerged for both a short and a longer adaptation duration, suggesting that these dynamics are robust across different exposure durations. This has practical implications for future paradigms investigating adaptation in the perception of causality: using shorter adaptors may suffice, allowing for faster trials without compromising the effects.

Our focus on recovery allows us to situate our findings within the spectrum of timescales of recovery observed in other adaptive processes. At the millisecond timescale, for instance, recovery from saccadic adaptation (e.g., Cassanello et al., 2016; Cassanello, Ostendorf, & Rolfs, 2019; McLaughlin, 1967; Rolfs, Knapen, & Cavanagh, 2010) and color adaptation (e.g., Zaidi et al., 2012) reveals aspects of recalibration that remain invisible when focusing solely on the adaptation phase (for two exceptions, see Ethier, Zee, & Shadmehr, 2008; Wong & Shelhamer, 2011). In contrast, some adaptive phenomena recover over much longer timescales: the McCollough effect can persist for days or even months (Jones & Holding, 1975), and chromatic shifts after donning tinted lenses fade gradually, documented as early as La Hire (1694/

2001). Positioned in the middle of this spectrum, our findings highlight how recovery trajectories differ across domains and show how directly studying recovery can improve our understanding of visual processing more broadly.

To our surprise, recovery profiles were inconsistent between the three experiments, varying both in their onset and whether they completely returned to baseline. In Experiment 1, recovery unfolded gradually but remained incomplete. In Experiment 2, recovery was likewise gradual, yet reached baseline levels, indicating complete recovery. In contrast, Experiment 3 showed an instantaneous onset of recovery following adaptation offset, but this recovery again remained incomplete. One possibility is that these differences reflect the coexistence of multiple adaptive processes operating on different timescales, including a rapidly reversible component alongside a more persistent, slowly adjusting component. Such a combination could potentially account for the rapid initial recovery observed in Experiment 3 while some adaptation nevertheless persisted following the 10-minute break. The persistence of residual adaptation despite the absence of structured visual input is broadly consistent with storage effects described in classical adaptation phenomena, such as the motion aftereffect, where adaptation can persist during periods without structured visual input (e.g., Mather, Verstraten, & Anstis, 1998). The present experiments do not allow us to distinguish these components directly. However, the variability across experiments suggests that recovery may reflect the interaction of multiple processes rather than a single uniform mechanism. A possible explanation for this variability is learning or memory consolidation across the experimental session: repeated exposure to launch events may stabilize adapted priors, thereby limiting a full recovery even when the adaptor is removed. Additionally, recovery dynamics may vary across individuals. If individuals differ in the speed with which they return toward baseline, the recovery process may not be fully captured within the current paradigm.

Considering these different potential explanations, we therefore do not interpret the present experiments as providing evidence for a single recovery mechanism although qualitatively all experiments suggest a rather fast recovery that reaches levels close to the pre-adaptation baseline. Rather, they demonstrate that recovery from adaptation can exhibit different temporal profiles under closely related experimental conditions, highlighting the need for future work to determine the processes underlying these differences.

4.3. *Mechanisms and methodological considerations in causal adaptation*

Another possibility is that the observed recovery goes beyond changes in sensory adaptation and also reflects additional influences of recent perceptual history. Previous work has shown that causality judgments are shaped by prior expectations and recent perceptual experience (e.g., Deodato & Melcher, 2022; Marsicano, Deodato, & Melcher, 2026), raising the possibility that repeated exposure to ambiguous post-adaptation stimuli gradually shifts observers' internal priors or decisional tendencies back toward baseline. More generally, attractive history-dependent biases and repulsive adaptation effects may coexist and interact during causal-event perception. The present experiments were not designed to disentangle these potentially interacting processes, and future work will be needed to determine how sensory adaptation, prior expectations, attention, and decisional factors jointly contribute to both the magnitude and recovery of adaptation effects.

A natural question is whether the recovery profile from adaptation would have looked different if observers had been deprived of *all* visual input between the adaptation and post-adaptation blocks. One way to achieve this would have been to ask observers to close their eyes during the break. Such a manipulation would provide a more stringent control than our podcast break, during which no changing visual stimuli were presented but observers still viewed a neutral display. However, any attempt to isolate the effect of elapsed time necessarily introduces additional factors that may influence recovery. For example, sitting with closed eyes in a dark room would not only reduce visual input but would also create a qualitatively different adaptation context. More generally, it is difficult to design a completely confound-free manipulation of elapsed time. These considerations highlight a broader methodological challenge: manipulations intended to isolate the contribution of elapsed time will inevitably alter other aspects of observers' perceptual state. Researchers designing such experiments should therefore carefully consider which aspects of the intervening period are being manipulated and which potential confounds are introduced. Given that we already observed an instantaneous recovery following the podcast break, it seems improbable that a slightly more extreme form of blocking visual input would produce a qualitatively different recovery profile.

5. Conclusion

Visual adaptation proved to be an effective and robust tool to address how spatially specific adaptation to causality is, and how the perception of causality recovers from adaptation. We provide compelling evidence that the mechanisms for processing causal interactions are highly spatially localized, rapidly responding to the onset of adaptation and then quickly recovering even without task-relevant input, eventually yielding nearly complete recovery that can feature slight differences to the pre-adaptation level. The combination of pronounced spatial specificity and instantaneous recovery supports our idea that the interpretation of causal interactions might originate at a visual, perceptual stage of processing.

CRediT authorship contribution statement

Laura van Zantwijk: Writing – original draft, Visualization, Software, Project administration, Methodology, Formal analysis, Conceptualization. **Martin Rolfs:** Writing – review & editing, Methodology. **Sven Ohl:** Writing – review & editing, Supervision, Software, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Ethics declaration

Written informed consent to take part in the study and to publish the article has been obtained from all participants or their legal representatives. The privacy rights of participants have been observed.

This study was performed in compliance with relevant laws, regulatory frameworks and guidelines where the research took place. This study was conducted in accordance with the Declaration of Helsinki. This study was approved by the Institut für Psychologie, Humboldt-Universität zu Berlin. (Approval No. 2022–58)

Declaration of competing interest

The authors declare no competing financial interests.

Acknowledgements

This research was supported by a DFG research grant to Sven Ohl (OH 274/4-1) as well as funding from the Heisenberg Programme of the DFG to Sven Ohl (OH 274/5-1).

Data availability

I have shared the link to my OSF repository.
Experiments 1, 2 and 3 (Original data and code) (OSF)

References

Addams, R. (1834). LI. An account of a peculiar optical phænomenon seen after having looked at a moving body. *The London, Edinburgh, and Dublin Philosophical Magazine*

and Journal of Science, 5(29), 373–374. https://doi.org/10.1080/14786443408648481
Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*(4), 433–436. https://doi.org/10.1163/156856897X00357
Cassanello, C., Ostendorf, F., & Rolfs, M. (2019). Differences in learning between global and vector-specific saccadic adaptation [conference abstract]. *Journal of Eye Movement Research, 12*(7), Article 168. https://www.eyemovement.org/pdf/ecem2019_abstracts.pdf.
Cassanello, C. R., Ohl, S., & Rolfs, M. (2016). Saccadic adaptation to a systematically varying disturbance. *Journal of Neurophysiology, 116*(2), 336–350. https://doi.org/10.1152/jn.00206.2016
Clarke, S., & Yousif, S. R. (2026). Can we "see" value? Spatiotopic "visual" adaptation to an imperceptible dimension. *Cognition, 266*, Article 106291. https://doi.org/10.1016/j.cognition.2025.106291
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
Deodato, M., & Melcher, D. (2022). The effect of perceptual history on the interpretation of causality. *Journal of Vision, 22*(11), 13. https://doi.org/10.1167/jov.22.11.13
Ethier, V., Zee, D. S., & Shadmehr, R. (2008). Spontaneous recovery of motor memory during saccade adaptation. *Journal of Neurophysiology, 99*(5), 2577–2583. https://doi.org/10.1152/jn.00015.2008
Freeman, J., & Simoncelli, E. P. (2011). Metamers of the ventral stream. *Nature Neuroscience, 14*(9), 1195–1201. https://doi.org/10.1038/nn.2889
Jones, P. D., & Holding, D. H. (1975). Extremely long-term persistence of the McCollough effect. *Journal of Experimental Psychology: Human Perception and Performance, 1*(4), 323–327. https://doi.org/10.1037/0096-1523.1.4.323
Kleiner, M., Brainard, D., Pelli, D., & Ingling, A. (2007). What's new in Psychtoolbox-3. *Perception, 36*, 1–16. https://doi.org/10.1068/v070821
Kohn, A. (2007). Visual adaptation: physiology, mechanisms, and functional benefits. *J. Neurophysiol., 97*.(5), 3155–3164. https://doi.org/10.1152/jn.00086.2007.
Kominsky, J. F., & Scholl, B. J. (2020). Retinotopic adaptation reveals distinct categories of causal perception. *Cognition, 203*, Article 104339. https://doi.org/10.1016/j.cognition.2020.104339
Kominsky, J. F., & Wenig, K. (2025). Causal perception(s). *Cognitive Science, 49*(9), Article e70107. https://doi.org/10.1111/cogs.70107
Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience, 10*(6), 779–786. https://doi.org/10.1038/nn1901
La Hire, P.de. (1694/2001). Les mémoires de mathématique et de physique tirez des registres de l'Académie royale des sciences. *De Diversis Artibus* (pp. 45–54). Brepols, ISBN 9782503511979. https://doi.org/10.1484/m.dda-eb.4.2017005. https://www.brepolsonline.net/content/books/10.1484/M.DDA-EB.4.2017005
Linares, D., & Lopez-Moliner, J. (2016). quickpsy: An R package to fit psychometric functions for multiple groups. *The R Journal, 8*, 122–131. https://doi.org/10.32614/RJ-2016-008
Marsicano, G., Deodato, M., & Melcher, D. (2026). Atypical integration of temporal evidence and priors in causality judgment along the autism-schizotypy continuum. *iScience, 29*(4). https://doi.org/10.1016/j.isci.2026.115325
Masland, R. H. (1969). Visual motion perception: Experimental modification. *Science, 165*(3895), 819–821. https://doi.org/10.1126/science.165.3895.819
Mather, G., Verstraten, F., & Anstis, S. M. (Eds.). (1998). *The motion aftereffect: A modern perspective*. MIT Press. https://mitpress.mit.edu/9780262133432/the-motion-after effect/.
McLaughlin, S. C. (1967). Parametric adjustment in saccadic eye movements. *Perception & Psychophysics, 2*(8), 359–362. https://doi.org/10.3758/BF03210071
Michotte, A. (1963/2017). *The perception of causality*. Routledge. https://doi.org/10.4324/9781315519050
Morey, R. D. (2008). Confidence intervals from normalized data: a correction to Cousineau (2005). *Tutorials Quant. Methods Psychol, 4*, 61–64. https://doi.org/10.20982/tqmp.04.2.p061
O'Neill, L. (2021, November 11). The science of birth without pregnancy [audio podcast]. https://podcasts.musixmatch.com/podcast/show-me-the-science-with-luke-oneill-01j3xaj4kfez2kfmq932p2peke/episode/49-the-science-of-birth-without-pregnancy-01j4aafqdkz7243m8t5ytqhd2.
O'Neill, L., 2024, November 7. The science of designer babies [audio podcast]. 189, https://podcasts.musixmatch.com/podcast/show-me-the-science-with-luke-oneill-01j3xaj4kfez2kfmq932p2peke/episode/189-the-science-of-designer-babies-01jcmtznrgm43v5jhrewmzvmse.
O'Neill, L., 2024, November 28. The science of heart attacks [audio podcast]. 192, https://podcasts.musixmatch.com/podcast/show-me-the-science-with-luke-oneill-01j3xaj4kfez2kfmq932p2peke/episode/192-the-science-of-heart-attacks-01jdvmxwrj8w3vsgpkhph5kbsy.
Ohl, S., & Rolfs, M. (2025). Visual routines for detecting causal interactions are tuned to motion direction. *eLife, 13*, Article RP93454. https://doi.org/10.7554/eLife.93454
Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*(4), 437–442. https://doi.org/10.1163/156856897x00366
Rips, L. J. (2011). Causation from perception. *Perspectives on Psychological Science, 6*(1), 77–97. https://doi.org/10.1177/1745691610393525
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
Rolfs, M., Knapen, T., & Cavanagh, P. (2010). Global saccadic adaptation. *Vision Research, 50*(18), 1882–1890. https://doi.org/10.1016/j.visres.2010.06.010
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
Sommer, B., Rolfs, M., & Ohl, S. (2025). Putting causality into context: Causal capture escapes the visual adaptation of causality. *bioRxiv*. https://doi.org/10.1101/2025.04.10.648104
Tregillus, K. E., & Engel, S. A. (2019). Long-term adaptation to color. *Current Opinion in Behavioral Sciences, Visual Perception, 30*, 116–121. https://doi.org/10.1016/j.cobeha.2019.07.005
Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron, 56*(2), 366–383. https://doi.org/10.1016/j.neuron.2007.10.012
Webster, M. A. (2015). Visual adaptation. *Annual Review of Vision Science, 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). In *The Oxford Handbook of Causal Reasoning* (pp. 1–32). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199399550.013.17.
Wong, A. L., & Shelhamer, M. (2011). Saccade adaptation improves in response to a gradually introduced stimulus perturbation. *Neuroscience Letters, 500*(3), 207–211. https://doi.org/10.1016/j.neulet.2011.06.039
Zaidi, Q., Ennis, R., Cao, D., & Lee, B. (2012). Neural locus of color afterimages. *Current Biology, 22*(3), 220–224. https://doi.org/10.1016/j.cub.2011.12.021