COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 49 (2025) e70107

ISSN: 1551-6709 online
DOI: 10.1111/cogs.70107

Causal Perception(s)

Jonathan F. Kominsky, Katharina Wenig

Department of Cognitive Science, Central European University

Received 5 March 2025; received in revised form 30 May 2025; accepted 31 July 2025

Abstract

In addition to detecting "low-level" features like shape, color, and movement, the human visual system perceives certain "higher-level" properties of the environment, like cause-and-effect interactions. The strongest evidence that we have true causal perception and not just inference comes from the phenomenon of retinotopically specific visual adaptation to launching, which shows that launching events have specialized processing at a point in the visual system that still uses the surface of the retina as its frame of reference. Using this paradigm, we show that the visual system adapts to two distinct causal features found in different types of interaction: a broad "launching-like" causality that is found in many billiard-ball-like collision events including "tool-effect" displays, "bursting," and event "state change" events; and an "entraining" causality in events where one object contacts and then moves together with another. Notably, adaptation to entraining is not based on continuous motion alone, as the movement of a single object does not generate the adaptation effect. These results not only

All procedures were reviewed and approved by the CEU Psychology Research Ethics Board (PREBO) under protocols 2022-17 and 2023-23. Experiment 1: Hypotheses, methods, and analysis plan were preregistered (https://osf.io/8tdq9/) on November 18, 2022 and updated (https://osf.io/g3kd5/) on November 26, 2022 following the discovery of a bug in the stimulus presentation code after a pilot participant run. Data collection started on November 27, 2022. All data and analysis scripts are publicly available at https://osf.io/38sd4/. Experiment 2: Hypotheses, methods, and analysis plan were preregistered (https://osf.io/kd5hx/) on May 25, 2023, and data collection commenced on May 30, 2023. All data and analysis scripts are publicly available at https://osf.io/38sd4/. Experiment 3: Hypotheses, methods, and analysis plan were preregistered (https://osf.io/hwdvt/) on September 4, 2023, and data collection commenced on September 5, 2023. All data and analysis scripts are publicly available at https://osf.io/38sd4/.

Correspondence should be sent to Jonathan F. Kominsky, Department of Cognitive Science, Central European University, Quellenstraße 51-55, 1100, Vienna, Austria. E-mail: Kominskyj@ceu.edu

demonstrate the existence of multiple causal perceptions, but also begin to characterize the precise features that define these different causal event categories in perceptual processing.

Keywords: Causal perception; Visual adaptation; Intuitive physics; Event representation

1. Introduction

When we see an event like the one rendered schematically in Fig. 1a (available in animated form at https://osf.io/2j937), we do not experience it as two independent movements, but instead automatically and irresistibly perceive that the left object (henceforth A) *causes* the the right object (henceforth B) to move (Hubbard, 2013; Michotte, 1963; Scholl & Tremoulet, 2000). This "launching" event has been studied extensively since Michotte's original work about it in the mid-20th century (Michotte, 1963), and the evidence is now very strong that this impression is not the product of some higher-level judgment or inference (Rips, 2011), but a true perceptual process (Hafri & Firestone, 2021; Scholl & Tremoulet, 2000): These events are distinguished from noncausal events (in which there is a long pause at the moment of contact, or A and B do not make contact at all) prior to conscious awareness (Moors, Wagemans, & de Wit, 2017), are only influenced by specific contextual information within a narrow spatial and temporal (±90ms) window (Choi & Scholl, 2006; Scholl & Nakayama, 2002), and are reliably detected as causal by human infants at 6 months of age (Cohen & Amsel, 1998; Leslie & Keeble, 1987; Saxe & Carey, 2006).

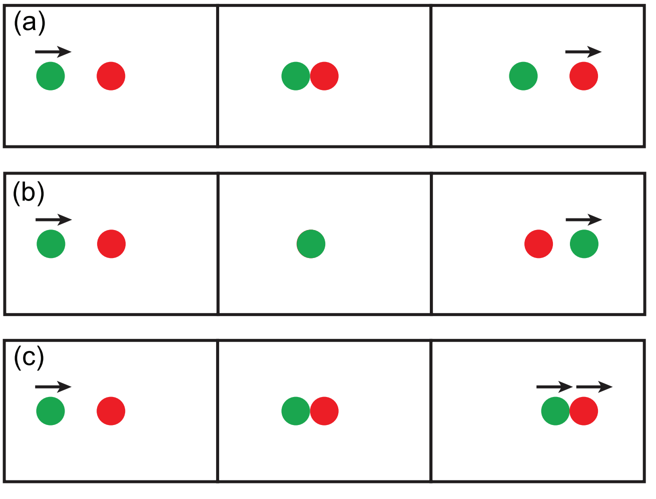

Fig. 1. Schematic depictions of three different events used in these experiments. (a) A protoyipcal "Launching" event, in which the green object appears to cause the green object to move. (b) A prototypical "Passing" event in which the green object appears to pass over the red one. (c) A prototypical "Entraining" event in which the green object collides with the red one and then they continue moving together.

One of the strongest pieces of evidence that launching causality is truly *perceived* is that it is subject to retinotopically specific visual adaptation (Rolfs, Dambacher, and Cavanagh, 2013). Visual adaptation effects are phenomena in which extended exposure to a particular stimulus causes subsequent ambiguous stimuli to appear less like the adapted stimulus. A classic example is the “waterfall illusion,” in which staring at constant downward motion (like a waterfall) for a long time and then looking at something that is stationary will cause illusory upward motion (Addams, 1834). An adaptation effect provides clear evidence that there is something in the visual system, which is specifically tuned to the adapted feature (Webster, 2016). Retinotopically specific adaptation effects are a subset of phenomena in which the adaptation effect is found exclusively (or mostly) when the adaptation stimulus and test stimulus are presented to the same location on the retina (e.g., color afterimages). Retinotopic specificity provides further evidence that the feature in question has “informationally encapsulated” perceptual processing (Fodor, 1983), in that, it is not only immune to explicit reasoning and judgment, but even to other information in the visual system that just happened to enter from a different location on the retina. Notably, retinotopically specific adaptation does not imply that the processing occurs on the retina itself; neuroimaging has found that the spatial structure of the retina is detectable in visual processing throughout much of the early visual system up to area V5/MT (Kolster, Peeters, & Orban, 2010).

In studies that apply this phenomenon to causality, participants first see a number of ambiguous events that can either be seen as launching or as noncausal “passing” (Fig. 1b), in which one object appears to pass over, under, or through the other. One can create events that are ambiguously one or the other by varying the degree to which A overlaps with B when A stops and B starts moving (Fig. 2a). After getting a baseline psychophysical function for the relationship between degree of overlap and the likelihood that participants report seeing passing, participants then see an adaptation stream of hundreds of unambiguous launching events. Following this adaptation, they are shown more ambiguous events. Multiple studies have found that, following this adaptation, participants reliably report seeing more passing events, that is, they show a visual adaptation effect (Karaminis et al., 2015; Ohl and Rolfs, 2025; Rolfs et al., 2013), and those that have tested for retinotopic specificity find that the effect is stronger when the test events are presented to the same retinotopic location as the adaptation stream (Kominsky & Scholl, 2020; Rolfs et al., 2013).

1.1. *Adapting “causality”?*

Adaptation effects depend on the presence of a particular feature or set of features within a stimulus, but any stimulus that includes that feature will generate the same adaptation effect (Webster, 2016). This is called “adaptation transfer.” Conversely, if a given stimulus lacks the key feature(s), it will not generate an adaptation effect.

If the launch/pass adaptation paradigm is truly demonstrating adaptation to “causality,” then any of the roughly dozen events that have been studied as “causal perception” (Hubbard, 2013) should generate an adaptation transfer effect on these launch/pass displays, while those described as noncausal should not. Past work has found that adaptation transfer to these launch/pass displays from events that vary the relative speed of the objects in the

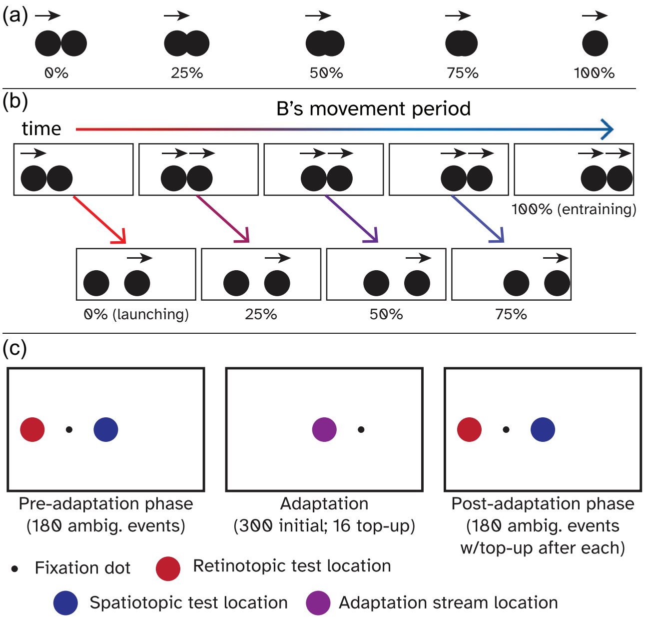

Fig. 2. Schematic depictions of the basic paradigm used in these experiments. (a) The “Launch/Pass” distinction used in Experiment 1. Participants saw launching-like events where the two objects overlapped to one of nine different degrees of overlap (here, we show only a subset). Zero percent overlap is prototypical launching (see Fig. 1a), while 100% overlap is prototypical passing (see Fig. 1b). (b) The “Launch/Push” distinction used in Experiments 2 and 3, where the two objects travel together for some percentage of B’s movement period, in nine steps (here, we show only a subset). If they travel together for 0% of the time, it is prototypical launching, while if they travel together for 100% of the time, it is prototypical entraining (see Fig. 1c). (c) Test events in the pre-adaptation block were presented at one of two locations (in red and blue). For the adaptation sequence, the fixation dot moved to the other side of the screen, and all adaptation events were presented at a single location (in purple). For the post-adaptation test block, the fixation dot moved back to the original side of the screen, and test events could now appear at the same retinotopic (in red) or spatiotopic (in blue) location as the test events.

event (described as “triggering” events; see Kominsky et al., 2017; Natsoulas, 1961) changes the colors of the agent and patient objects between the adaptation and test stimuli (Ohl & Rolfs, 2025), and an event where A slows but does not fully stop following contact with B (Kominsky and Scholl, 2020). Critically, there is evidence that the adaptation effect only applies to *causal* events: adaptation to noncausal “slip” events, in which A moves through B and stops on the far side before B starts moving, does not generate an adaptation effect (Rolfs et al., 2013).

However, the causal events that have demonstrated these adaptation or adaptation transfer effects are all parametric variations on launching. With the exception of triggering, they are all described as “launching,” and even triggering only differs from launching in the ratio of relative speeds of the objects involved exceeding a 1:2 threshold (Michotte, 1963; Natsoulas, 1961). In short, the demonstrations of the adaptation effect to date have not covered a broad enough set of events to make the claim that the feature being adapted is an abstract property of “causality.”

Moreover, Kominsky and Scholl (2020) showed that there is one event which has been described as another prototypical example of "causal perception" but which does *not* generate this adaptation transfer effect: causal "entraining" events (Fig. 1c; Michotte, 1963; see also https://osf.io/m2js9). Entraining events start similarly to launching events, but when A contacts B, the two objects remain in contact and move together. From this failure of adaptation transfer alone, it is clear that the feature being adapted cannot be an abstract feature of "causality" writ large.

One obvious concern is that the failure of adaptation transfer with entraining means that the other adaptation effects can be explained through some feature of these events that is unrelated to causality. The feature that entraining removes from launching is whether A stops, which at first glance would seem to be a clear candidate for the feature being adapted. However, there is already clear evidence that "A stopping" is neither necessary nor sufficient to drive these adaptation effects: Noncausal "slip" events include A stopping (albeit on the far side of B) but do not generate an adaptation transfer effect (Rolfs et al., 2013), and adaptation to a launching event in which A continues moving after contact but at a slower speed still generates a retinotopically specific adaptation transfer effect (Kominsky and Scholl, 2020).

There is one other reported case of failed transfer, in cases when participants are adapted to launching all in one direction and then tested on launch/pass events on an axis of motion that is more than 30 degrees off of the one they were adapted to (Ohl & Rolfs, 2025). While this suggests that the adaptation effect is somewhat direction-dependent, every other demonstration of the effect adapted and tested participants with events on random axes of motion, suggesting that it is not a traditional directional motion adaptation effect. If it were, one would expect that adaptation to events with random directions of motion would simply cancel out, but this is not the case. So, some consistency in direction of motion is necessary between adaptation and test, but a motion aftereffect is not sufficient to fully explain the effect. However, this study also did not test for retinotopic specificity, so it is unknown whether the direction sensitivity is retinotopically specific or not.

To summarize what we know from previous work using the launch/pass adaptation paradigm: Matching speeds, either within the event or between adaptation and test, are not necessary. Consistent object features or associations of features with role (agent/patient) is not necessary. Object A stopping is not necessary, nor is it sufficient (in the case of the "slip" event). Perhaps most critically, "causality," at least defined in terms of apparent contact, is not sufficient, but may be necessary.[1]

This complex pattern of results leaves two glaring questions: First, exactly what feature or features are being adapted in launching adaptation? Second, is entraining actually a distinct causal *percept*, or merely an event that is *described* as causal on the basis of some more abstract inference, and not actually a privileged category in perception at all?

1.2. The features of causal events

In the literature, there are several other events which are still described as causal, but that lack specific features of the launching event. For example, "tool effect" events in which B is at rest in contact with an object C, and A contacts C, causing B to move away from C (think of a "Newton's cradle" toy, or see animation at https://osf.io/hzvwk/). This event, while described as A causing B to move (Hubbard & Favretto, 2003) and even understood as causal by 8-month-old infants (Cohen, Rundell, Spellman, & Cashon, 1999), does not involve direct contact between A and B. Other events change the nature of the effect, that is, what happens to B. In “bursting” or “shattering” events, B breaks into many smaller pieces that move away from the point of contact with A (Hubbard & Ruppel, 2013; White & Milne, 1999; animation: https://osf.io/wxajm), while in “state change” events, B does not move at all, but instead changes color or shape (Adibpour & Hochmann, 2023; Liu, Brooks, & Spelke, 2019; Muentener & Carey, 2010).

These events do not just differ from launching parametrically along a single dimension, they differ in the features which are actually present in the event (direct contact, object cohesion, and B’s movement, respectively). Nonetheless, they are still described as causal (Hubbard, 2013; see also Experiment A1 in the Supplemenatry Appendix). If adapting to one of these events failed to produce an adaptation transfer effect, it would be clear evidence that the feature that particular event lacks is necessary for the adaptation effect. However, if all of these events produced adaptation transfer, it might suggest that the adaptation effect is not reducible to a single necessary and sufficient feature, but rather to a combination of features that is found in several launching-like causal events.

As for entraining, the failure of adaptation transfer could either indicate that entraining is not truly *perceived* as causal at all and only described as such due to some more domain-general learning mechanism (e.g., Benton, 2024; Rips, 2011; Ullman, Harari, & Dorfman, 2012), or that entraining is its own distinct category of causal event that still has specialized perceptual processing and the launch/pass distinction was simply the wrong feature dimension to use as a test. Just because an adaptation effect is not found with one test stimulus does not necessarily mean that a given feature or event is not subject to adaptation. For example, you get different motion aftereffects depending on whether you use static or dynamic test stimuli, and second-order motion aftereffects are only detectable with static test stimuli (Mather, Pavan, Campana, & Casco, 2008). Importantly, Kominsky and Scholl (2020) could not conclude that entraining was not subject to adaptation, just that it did not adapt the launch/pass distinction.

The feature that distinguishes entraining from launching is whether A continues to move with B after contact. Notably, this feature can be manipulated along a continuous spectrum: Does A move with B for the full duration of B’s motion, or does it stop at some point? Cases in which A moves with B for part, but not all, of its motion are sometimes described as “launching by expulsion” (Michotte, 1963; see animation: https://osf.io/j6v3y), but since this feature dimension distinguishes launching from entraining, it may be possible to find a point along this dimension that is ambiguous between launching and entraining. If so, it provides an opportunity to test whether adapting to launching creates an adaptation effect along this feature dimension as well, but more critically, to see if entraining generates an *opposed* adaptation effect along this feature dimension.

1.3. The current experiments

There are two goals for the current experiments. The first goal is to better characterize exactly what features perception is sensitive to that drives these visual adaptation effects for launch/pass displays. To this end, Experiment 1 tests adaptation transfer using the same launch/pass displays as previous studies with three new adaptation events that each remove a single feature of launching, but are still described as causal in explicit reports.

The second goal is to determine whether entraining is a distinct perceptual category, that is, positive evidence for multiple causal perception*s* by characterizing the features that make up a second category. Experiments 2 and 3 ask whether entraining is a distinct, or even opposed, category of causal percept by testing whether it generates adaptation effects for a novel test stimulus that we describe as "launch/push." In particular, Experiment 2 tested whether adaptation to launching versus entraining would generate *opposed* adaptation effects on this feature dimension, which would provide initial evidence whether entraining is a distinct category of event that is subject to visual adaptation, or merely "not-launching" (e.g., like "slip" events). Experiment 3 then tests whether adaptation to entraining can be explained by the continuous movement of object A, or if it requires an actual *causal* entraining event.

2. Experiment 1

We first set out to better characterize the feature driving the launch/pass adaptation effect by testing adaptation transfer from three events which are all described by adults as "causal" (Hubbard, 2013; see also Experiment A1 in the Supplemenatry Appendix), but each of which removes one feature of the launching event (see Fig. 3b–d and Movies S2–S5, found at https://osf.io/38sd4/files/osfstorage under "Supplementary videos").

2.1. Methods

All experiments were preregistered. All preregistrations, experiment code, anonymized data, and analysis code and outputs can be found at https://osf.io/38sd4/. All study procedures were approved prior to data collection by the CEU Psychology Research Ethics Board (PREBO) as part of the protocol 2022-17, "Behavioral studies of causal reasoning and perception," and all participants provided written informed consent prior to participation.

2.1.1. Participants

Based on Experiment 2 of Kominsky and Scholl (2020), which was the closest methodological analog to the current experiments, we were confident that adaptation transfer effects, even for untested events, could be reliably detected with 20 participants in a single session. We determined that a sample size of 20 would provide 80% power to detect an adaptation effect in a single condition as measured by a two-tailed single-sample t-test with an effect size of $d = .66$, which is smaller than most of the significant effect sizes observed in Kominsky and Scholl (2020)'s Experiment 2. We, therefore, used a target sample size of 20 participants per adaptation condition for all of the experiments reported here.

In Experiment 1, we aimed to recruit 20 participants in each between-subjects condition. All participants were recruited from the CEU Cognitive Science Department SONA systems database, which includes members of the CEU community as well as the general population.

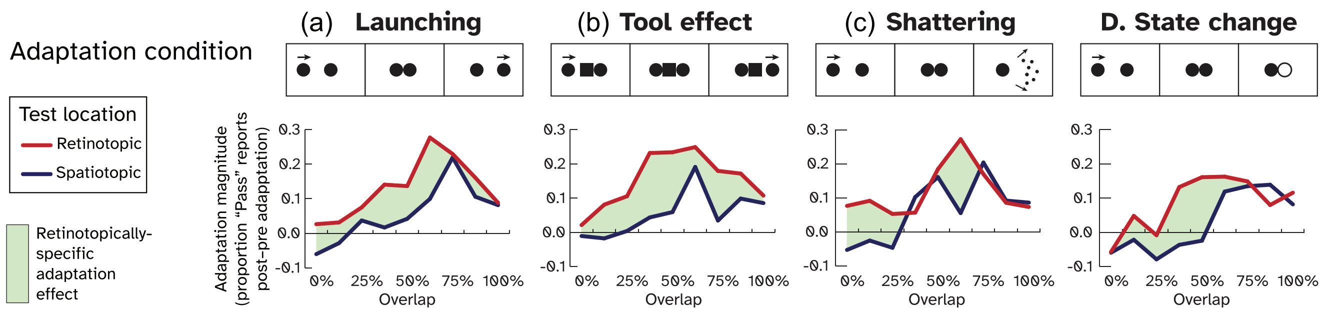

Fig. 3. Conditions and results of Experiment 1. Top row shows a schematic version of the adaptation event (videos available at https://osf.io/38sd4). Each graph shows the magnitude of the adaptation effect, operationalized as the proportion of "pass" reports at each degree of overlap and each location in the post-adaptation block minus the proportion of "pass" reports at the corresponding overlap and location in the pre-adaptation block. Lines show average adaptation magnitudes at the two test locations (retinotopic in red and spatiotopic in blue, see Fig. 2c), and the shaded area in green indicates the average magnitude of *retinotopically specific* visual adaptation.

For Experiment 1, we recruited 80 participants (21 male, 56 female, 3 nonbinary or other). An additional 34 participants were recruited but excluded for failing to pass our preregistered exclusion criteria (see below), either prior to the adaptation block or prior to analyses, and replaced until we hit the target sample size. All participants were compensated at a rate of €10/h for a study that typically took about an hour.

2.1.2. *Apparatus*

The experiments were all programmed in Python using the PsychoPy library (Peirce et al., 2019). We used a GazePoint GP3 eye-tracker operating at 60 Hz to ensure participants fixated on a fixation point for the purposes of testing retinotopic specificity. Participants sat in a chair and rested their head in a chin-rest 70 cm from a 41 cm by 72 cm 60 Hz LCD monitor operating at 1280x720 resolution.

2.1.3. *Stimuli and procedure*

After filling out an informed consent form, participants completed an eye-tracker calibration and then read instructions about the task, explaining that they would see events that looked like either a "Launch" in which one object collided with and caused the other to move, or a "Pass" in which one object passed over, under, or through the other without making contact. They were instructed to press the "L" key if the event looked like a launch, and the "P" key if it looked like a pass. They then saw three prototypical examples of each type of event (0% overlap events for launching, 100% overlap events for passing) at a reduced speed. They then completed 18 practice test trials, and for each one had to indicate whether they saw it as a "launch" or a "pass" by pressing the L or P keys, respectively. These practice trials were not included in any analyses.

Each test event consisted of a fixation dot subtending .5 degrees of visual angle (dva) and appearing 5 dva to the left or right of the center of the screen (side counterbalanced between subjects within each condition). Throughout the experiment, including this training, a test event was only displayed after the eye-tracker detected a fixation within 2.5 dva of the fixation dot. When the eye-tracker detected a fixation at the fixation dot, two circles appeared, each one 1.5 dva in diameter. One circle, object B, appeared 5 dva to the left or right of the fixation dot, while the other, object A, appeared between 5 and 6 diameters away from it in a random direction. Object A immediately started moving toward object B for 6 frames (at 60 frames per second) at a speed of 2 dva/frame, after which object A stopped and object B immediately began moving in the same direction and at the same speed for an equal number of frames. Varying the starting distance of object A meant that when A stopped and B started moving, the two objects overlapped to a variable degree from 0% (adjacent, prototypical launching) to 100% (full overlap, prototypical passing), in steps of 12.5%. This yielded nine distinct degrees of overlap. (See Movie S1 for examples.)

The pre-adaptation test phase consisted of 180 trials, 10 at each combination of location and overlap. For each one, participants watched an event and then pressed the "L" key if they saw it as a "launch" and the "P" key if they saw it as a pass. After a 300 ms inter-trial interval (ITI), the next test trial started as soon as the eye-tracker detected a valid fixation.

Using the same exclusion criteria from Kominsky and Scholl (2020), at the end of the pre-adaptation block, we excluded participants who could not make the relevant distinction at this speed of presentation. To be included, participants had to respond “launch” more than 50% of the time for the two lowest degrees of overlap/shared movement combined (0% and 12.5%), and less than 50% of the time for the two highest degrees of overlap/shared movement combined (87.5% and 100%), with at least 20% difference between the two (e.g., 40% vs. 60%). This was automatically computed by the presentation script, and if a participant did not meet these criteria, the script saved the data from the pre-adaptation block and quit. In this experiment, 30 participants were excluded on this basis (27.2% attrition), which is in line with previous work and most likely due to the high speed of the event (see Kominsky and Scholl, 2020 for further discussion of why the speed of these events leads to relatively high attrition; we continue to use these speeds because speed trades off with number of trials and length of testing session).

For participants who moved on to the adaptation block, the first thing they saw was the fixation dot moving from its location during the test events to a location 5 dva above the center of the screen. When the eye-tracker detected a fixation at this location, the dot then moved to a location 5 dva from the center of the screen on the opposite side from its location during the pre-adaptation block (see Fig. 2c, center panel). Participants then saw 400 adaptation events, all presented at the center of the screen. Participants were instructed not to respond to these adaptation events, and they proceeded automatically with a 100 ms ISI as long as the eye-tracker detected a fixation within 2.5 dva of the fixation dot. If participants’ fixation drifted from the dot, the adaptation sequence would pause until they looked at it again. This ensured that the adaptation events were presented to a consistent location on the retina.

There were four different adaptation conditions in this experiment (see the top row of Fig. 3 and Movies S2–S5), and participants were assigned randomly to one of them. In the “launching” adaptation condition, the adaptation event consisted of 0% overlap launching events identical to those used in the test block.

In the “tool effect” condition, the adaptation event was similar to launching, but with the addition of a square subtending 1.5 dva which started adjacent with B. Object A moved until it was adjacent with this square, at which point A stopped and B began moving. This event is described as causal (Michotte, 1991), but compared to launching, removes the feature of direct contact between A and B.

In the “shattering” condition, at the moment of contact, B was replaced by nine smaller circles, each one subtending .167 dva, which moved linearly in a random direction within ±45 degrees of the axis of A’s motion, but with the same speed as A. This event is described as causal (Hubbard & Ruppel, 2013; White & Milne, 1999), but compared to launching, removes the feature of B’s objecthood following contact.

In the “state change” condition, the adaptation event was identical to launching except that instead of moving at the moment of contact, B changed from black to white and remained stationary for 6 frames. This event is treated as causal by infants when the agent is a human hand (Liu et al., 2019; Muentener & Carey, 2010), but compared to launching, removes the feature of B's movement.

Following the initial adaptation stream, the fixation dot moved back to the other side of the screen and participants saw one test event. Then, the fixation dot once again returned to the opposite side of the screen and participants saw 16 "top-up" adaptation events. This was repeated for 180 trials. The test events were identical to the test events in the pre-adaptation block.

2.1.4. *Eye-tracking exclusions*

Participants' eyes were tracked at 60 Hz throughout the experiment, and test trials where participants made saccades during the animation (defined as movement between two samples greater than 5 SDs from the total median velocity for that participant across all test trials; Engbert & Kliegl, 2003) or the eye-tracker provided invalid samples (usually due to blinks) were excluded by a data processing function that ran automatically at the conclusion of the experiment while saving the data file. This ensured that we only analyzed responses to events that were actually viewed in the intended location (retinotopically matched or mismatched to adaptation location). At the end of the experiment, participants who did not have at least 240 valid test trials across both blocks, and at least one valid trial at each overlap for each location in each test block, were automatically flagged for exclusion by this data processing function, and their data was replaced by a later participant. We excluded four participants on this basis.

2.2. *Results*

The results can be found in Fig. 3, which shows the magnitude of the adaptation effect (proportion passing reports post-adaptation – proportion of passing reports pre-adaptation) for each location and degree of overlap in each condition. Our preregistered analysis plan (https://osf.io/38sd4) examines difference between test locations in adaptation magnitude (green shaded area) averaged across all degrees of overlap, following the logic outlined in Kominsky and Scholl (2020). An alternative analysis approach using Bayesian estimates of points of subjective equality (PSEs) following Rolfs et al. (2013) can be found in the "Supplemental Analyses" in Appendix B, and support the same conclusions as the analyses reported here. This analysis also includes the estimated PSE for each condition in Table B1.

All four conditions showed a significant retinotopically specific adaptation effect. Table 1 shows the average magnitude of the retinotopically specific adaptation effect and a two-tailed one-sample *t*-test of this value against 0. To determine if there were differences in the magnitude of retinotopically specific adaptation depending on the adaptation condition, we conducted a one-way ANOVA, which found no significant effect of condition, $F(3, 76) = 0.59$, $p = .62$.

Table 1
Magnitude of retinotopically specific adaptation effect in Experiment 1

Adaptation condition	*M* ret. spec. adaptation (*SD*)	*t*-test versus 0
Launching	.073 (.269)	$t(19) = 2.19$, $p = .041$, $d = .49$
Tool effect	.099 (.256)	$t(19) = 3.90$, $p < .001$, $d = .87$
Shattering	.054 (.256)	$t(19) = 2.16$, $p = .044$, $d = .48$
State change	.059 (.278)	$t(19) = 2.87$, $p = .01$, $d = .64$

Note. All values calculated as magnitude of retinotopic adaptation - magnitude of spatiotopic adaptation, then averaged across all degrees of overlap.

2.3. Discussion

Our results indicate that the launch/pass adaptation effect is not based on any one feature of launching, but more likely a combination of features shared by all of these events. Notably, all of these events are still described as causal (which we tested in Supplementary Experiment A1, and you can experience for yourself by watching the Movies S2–S5 at https://osf.io/38sd4/files/osfstorage), and indeed, these results suggest they may also be truly *perceived* as causal, not merely described as such. One salient question is whether the adaptation is simply to the feature of "A stopping." However, as discussed in the introduction, "A stopping" is neither necessary nor sufficient to produce an adaptation effect. In Kominsky and Scholl (2020)'s Experiment 3, adaptation to a launching event in which A does not stop but continues moving at a slower speed following contact generates an adaptation transfer effect. As for sufficiency, the noncausal "slip" event used by Rolfs et al. (2013) involves A stopping (after it has passed through B), but did not generate an adaptation effect. It may still be the case that A stopping *on initial contact with another object* is a sufficient feature to drive adaptation, but even then, one can argue that this is a causal interaction (albeit one that people rarely acknowledge in describing these events; Mayrhofer & Waldmann, 2014; White, 2009).

There are some more abstract features that are shared by all of these causal events, notably the assignment of "agent" and "patient" roles. Recent work has suggested that it is plausible that the visual system could be sensitive to this relationship. When human figures are involved, 6-month-old infants (Papeo et al., 2024) and adults (Vettori, Odin, Hochmann, & Papeo, 2025) automatically process agent and patient roles, even without seeing a direct interaction. Six-month-old infants also track agent and patient roles in causal launching events (Leslie & Keeble, 1987). However, Ohl & Rolfs (2025) found that adapting to a launching event in which a red object launched a green object transferred to a test event in which their causal roles were reversed. However, rather than adapting to a specific object in a specific role, it is possible that participants in the current experiment adapted to the assignment of roles in the event, that is, the fact that there was an unambiguous agent and unambiguous patient. Passing events have no such roles, as they do not involve an interaction between two objects, so if that is the basis of the ambiguity between launching and passing, adapting to something like "event role assignment" could be a reasonable explanation for these results. However, whether this is a more parsimonious explanation for these transfer effects than a combination of "lower-level" spatiotemporal features that are distinctive to causal interactions will depend on future work investigating whether event roles are encoded early in the visual system.

From our current results, we argue that the visual system has specialized processing for a combination of features that is particular to causal events. However, as Kominsky and Scholl (2020) showed, this combination of features is not found in *all* causal events. Indeed, even if we take the "event role" explanation mentioned above, we might expect a failure of transfer from entraining to the launch/pass distinction, as the causal roles in an entraining event become ambiguous after contact (it is indistinguishable whether A is pushing B or B is pulling A). This leaves us with a crucial question: does the visual system have distinct specialized processing for causal entraining?

3. Experiment 2

Whereas Experiment 1 used the same launch/pass test events as previous work, Experiments 2 and 3 used a novel "launch/push" contrast (Fig. 2b), which was designed to be ambiguous between launching events and entraining events. The launch/pass contrast is based on the feature of overlap, how much A overlaps B when A stops and B starts moving. The launch/push contrast is based on the feature of where in B's motion A stops moving (and B continues). In prototypical launching, A stops immediately when B starts moving. In prototypical entraining, A moves with B for all of B's motion. This feature, therefore, distinguishes launching from entraining, much as overlap distinguishes launching from passing. Assuming that there is a degree of overlap that is ambiguous between launching and entraining (which, to foreshadow our results, there is), then this feature dimension might also be subject to adaptation. In other words, if we adapt to launching and test on this launch/push display, people should report more "push" events (and this adaptation effect may or may not be retinotopically specific). More importantly, if entraining is a distinct causal percept that the visual system is sensitive to (and not merely "not-launching"), then adapting to entraining should also generate an adaptation effect *in the opposite direction*: Adapting to entraining should lead people to report more "launch" events.

3.1. Methods

3.1.1. Participants

We recruited 40 participants (11 male, 24 female, 5 did not report) who had not participated in Experiment 1 from the same population and from the Vienna Cognitive Science Hub Study Participant Platform (Bock, Baetge, & Nicklisch, 2014). Compensation was the same as Experiment 1.

3.1.2. Stimuli and procedure

The primary difference between Experiments 2 and 3 and Experiment 1 was in the feature dimension used to create the test events. Every event had 0% overlap, but test events in this experiment varied the proportion of B's movement during which A remained in contact with

B and they moved together (or put differently, when A stopped moving following contact with B). In each test event, object A always appeared 6 diameters away from object B, and the events were slowed down so that A's movement took 8 frames (at 60 fps), until it was fully adjacent with B, at which point, B moved for 8 frames. At the moment of contact, A would continue to move with B for 0–8 frames, creating nine different durations of shared contact (see Fig. 2b and Movie S6 for examples). If A stopped at frame 0, the event looked like prototypical launching. If it moved with B for all 8 frames, it looked like prototypical entraining. Participants saw three examples each of these prototypical events as part of the instructions (much as they saw prototypical launch and pass events in Experiment 1).

The size of the objects, the distance between the objects and the fixation dot, and the gaze-contingent criteria for starting a trial were all identical to Experiment 1. The instructions to participants were also similar except that they were told they would see an event that would either "look like one circle collides with and 'launches' the other, or one circle will appear to 'push' the other along, moving with it for a while and then releasing it." The same response keys (L and P) were used for this experiment.

There were only two adaptation conditions, "launching" and "entraining." The launching adaptation condition used the same event as the test events, with 0 frames of shared movement. The entraining condition also used the same event as the test event, with 100% shared movement. The adaptation procedure was identical to Experiment 1, as were the exclusion criteria. In this experiment, 16 participants were excluded for failing to adequately distinguish launching from pushing (28% attrition), which is comparable to the exclusion rate in Experiment 1. Three additional participants were excluded for failing to provide a sufficient number of valid trials.

3.2. Results

One question going into Experiment 2 was whether participants would even be able to make the "launch/push" distinction reliably, since it had never been used before. This was implicitly tested by our exclusion criterion, as it requires participants to provide mostly "launch" responses when the co-travel time is short and mostly "push" reports when it is long. In this experiment, 27/67 (40%) participants failed to meet this criterion, though this was not notably different from the rate of exclusions based on failing to make the launch/pass distinction in Experiment 1 (34/114; 30%), and to foreshadow, we found a much lower exclusion rate in Experiment 3 which used the same test events and exclusion criteria. As with previous work with untrained observers, the relatively high (and variable) exclusion rate is most likely due to the high speed of the displays (for further discussion of the speed issue, see Kominsky and Scholl, 2020).

As this was the first time this distinction has been used, we provide more granular graphs here showing the raw proportion of launch and pass reports at each degree of overlap for each block and location in Fig. 4.

We graph the magnitude of the adaptation effects in Fig. 5a,b (these are the same data as Fig. 4, but graphing the height of the difference between the pre- and post-adaptation curves). Note that an adaptation effect above the x-axis indicates more "push" reports following adap-

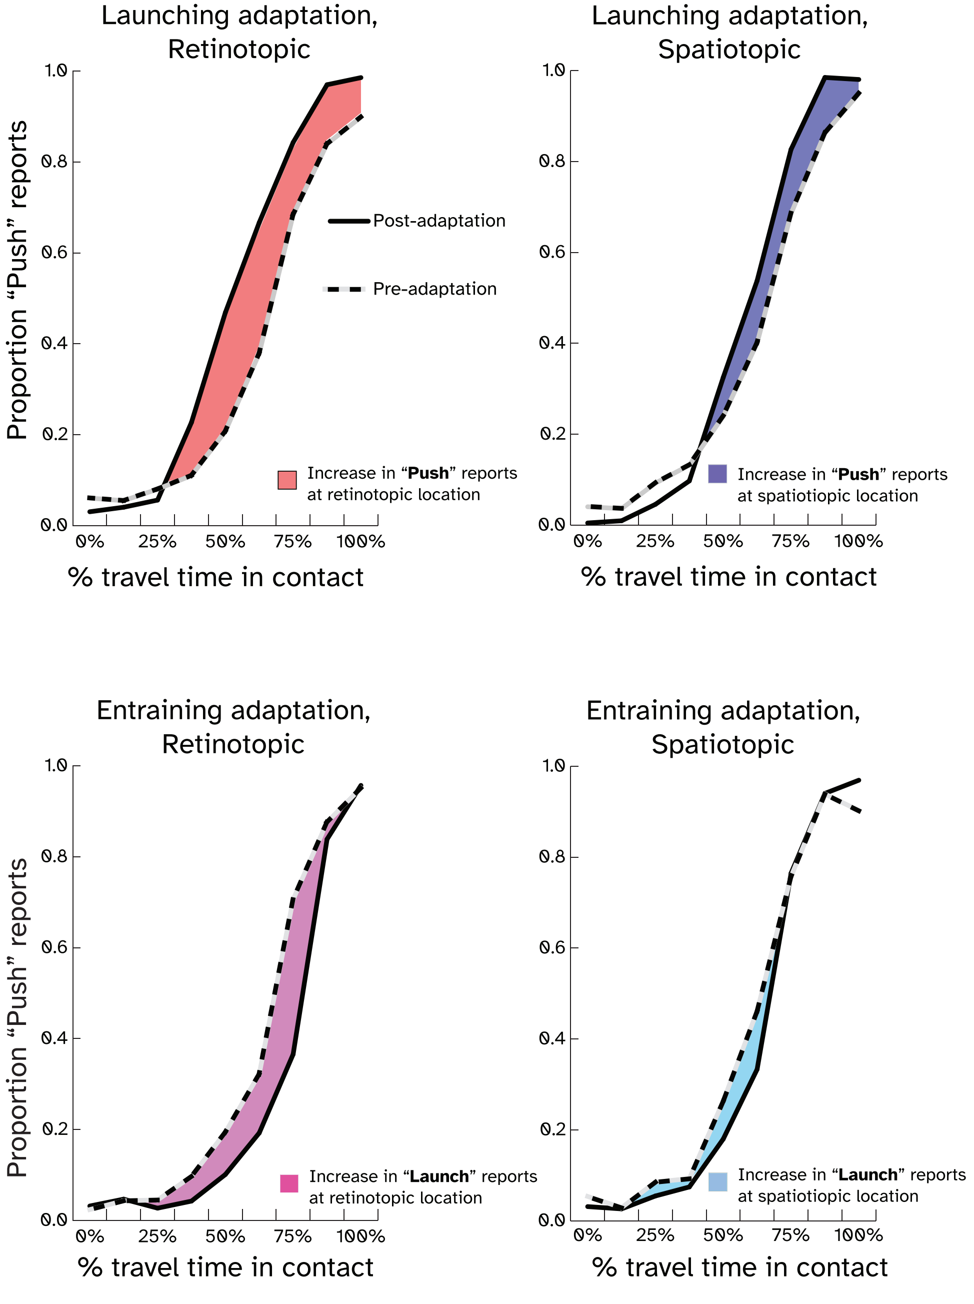

Fig. 4. Average "Push" reports in each condition and location in Experiment 2. Dashed lines show pre-adaptation responses, solid lines show post-adaptation responses. Note that the direction of the adaptation effect reverses in the entraining adaptation graphs, such that there are *fewer* "Push" reports post-adaptation.

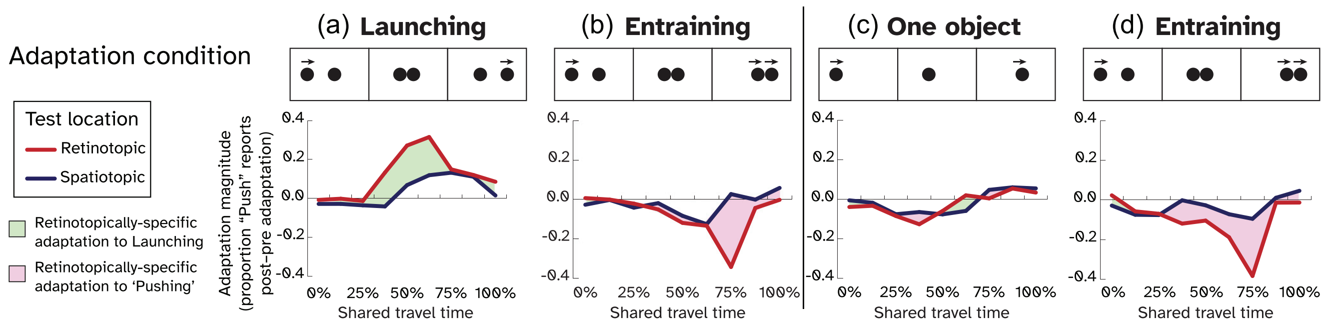

Fig. 5. Conditions and results of Experiments 2 and 3. Top row shows a schematic version of the adaptation event (videos available at https://osf.io/38sd4). Each graph shows the magnitude of the adaptation effect, operationalized as the proportion of "push" reports at each degree of overlap and each location in the post-adaptation block minus the proportion of "push" reports at the corresponding overlap and location in the pre-adaptation block. Lines show average adaptation magnitudes at the two test locations (retinotopic in red and spatiotopic in blue, see Fig. 2c). The shaded area in green indicates the average magnitude of retinotopically specific visual adaptation to launching (i.e., more "push" reports), while the shaded area in magenta indicates the magnitude of retinotopically specific adaptation to "pushing" (i.e., more "launch" reports).

tation, while an adaptation effect below the x-axis indicates more "launch" reports following adaptation, while a line around the x-axis indicates no adaptation effect at all.

As these figures suggest, we found the predicted opposed retinotopically specific adaptation effects. Our preregistered analysis consisted of a two-tailed *t*-test against 0 in each condition, and an independent-samples *t*-test comparing the two conditions. We found a significant retinotopically specific visual adaptation effect in the launching adaptation condition ($M = 0.082$, $SD = 0.236$), $t(19) = 3.63$, $p < .01$, $d = 0.812$, and critically, a significant retinotopically specific visual adaptation effect *in the opposite direction* in the entraining adaptation condition ($M = -0.055$, $SD = 0.228$), $t(19) = -3.36$, $p < .01$, $d = -0.75$. The conditions also differed significantly from each other, $t(38) = 4.91$, $p < .001$, $d = 1.55$. In short, there was significant retinotopically specific adaptation in both conditions, but in opposite directions. We also conducted a Bayesian PSE analysis for this experiment (see Supplemenatry Appendix), which produced qualitatively identical results.

3.3. Discussion

There are two important results of this experiment. The first is that retinotopically specific visual adaptation to launching is not restricted to the launch/pass distinction. This result alone indicates that components of perceptual processing that are affected by the adaptation stream are not just sensitive to one movement (passing) versus two movements (launching): in these launch/push displays, there are always two movements. In addition, it does not rely on ambiguity about object identity: in passing, one object appears to move across the screen, while the other remains stationary in the center, while here, the identity of each object is unambiguous throughout the events. Instead, these results reinforce the conclusion that the visual system is sensitive to a *combination* of features found in launching and the launching-like causal events studied in Experiment 1.

Second, and perhaps more importantly, we found evidence for adaptation to entraining events such that adaptation to entraining led to more *launching* reports. Furthermore, this effect was also retinotopically specific. The hypothesis we sought to test with the entraining adaptation condition was whether causal entraining has a distinct feature or set of features that could be adapted, and the results clearly show that it does. Thus, entraining may be a separate category of causal event in *perception*. Notably, while the adaptation effect is numerically smaller in absolute magnitude ($M = 0.082$ for launching, $M = -0.055$ for entraining), that is, primarily due to the fact that the pre-adaptation PSE for the launch/push distinction shows participants had a general bias toward launches. If you look at Fig. 4, you will notice that the dashed lines (representing the pre-adaptation responses) intersect .5 on the y-axis some ways past the 50% point on the x-axis. More precisely, Appendix C reports a Bayeisan PSE analysis and Table C1 reports the mean PSEs before and after adaptation; the mean PSE before any adaptation across conditions is around 63% shared travel time. This bias means that most of our test events were seen as "launching" to start with, and, therefore, there were fewer events that could be made to look more like "launching" by adapting to entraining. Thus, this asymmetry is most likely related to how many test events *could* show a strong entraining adaptation effect.

However, what this experiment did not fully establish was whether the adaptation effect for entraining was adaptation to the *causality* of entraining, so in Experiment 3, we compared adaptation to entraining to adaptation to a very similar noncausal event.

4. Experiment 3

The goal of this experiment was straightforward: Replicate the novel entraining adaptation effect from Experiment 2, and compare it to a noncausal entraining-like adaptation event to determine if the adaptation effect is to a feature of *causal* entraining events, or to a feature that is unrelated to causality. The most likely candidate feature of entraining that could drive this effect is the continuous movement of object A, especially since the launch/push distinction is based on whether and when A stops. Therefore, we adapted participants to an event that was identical to entraining but without an object B, in other words, just showing a single object moving continuously.

4.1. Methods

4.1.1. Participants

We recruited 40 new participants (11 male, 22 female, 2 nonbinary or other, 5 did not report) from the same sources as Experiment 2.

4.1.2. Stimuli and procedure

This experiment was identical to Experiment 2, except that the "launching" adaptation condition was replaced with a "one object" adaptation condition. The "one object" adaptation condition used an event that was identical to the entraining adaptation condition, but without object B. In other words, A just moved continuously for 16 frames, and no other objects were visible during the event. In this experiment, 12 participants failed to adequately distinguish launching from pushing, and two were excluded for failing to provide a sufficient number of valid trials (14/54, 26%).

4.2. Results

As suggested by Fig. 5d, we successfully replicated the significant adaptation effect in the entraining adaptation condition such that people reported seeing more launches following entraining adaptation ($M = -0.067$, $SD = 0.240$), $t(19) = -3.25$, $p < .01$, $d = -0.73$, but as seen in Fig. 5c, there was no such effect in the one-object adaptation condition ($M = -0.010$, $SD = 0.210$), $t(19) = -0.58$, $p = .57$, and there was a significant difference between the two conditions, $t(38) = -2.07$, $p = .046$, $d = -0.65$, indicating that the noncausal adaptation condition was truly different from entraining adaptation. A Bayesian PSE analysis (see Appendix D) produced qualitatively identical results, and replicated the finding of the previous experiment that the PSEs before any adaptation were over 50% (see Table D1).

5. General discussion

These three experiments characterize two distinct, and indeed opposed, categories of causal events in human perceptual processing: A "launching-like" event category that encompasses all of the adaptation events in Experiment 1 plus events with different object identities (Ohl & Rolfs, 2025), as well as "triggering" events (Kominsky and Scholl, 2020); and an "entraining" category that thus far consists of entraining alone (Experiments 2 and 3). Not only does this demonstrate that the visual system is not sensitive to "causality" in the most general sense, it goes beyond Kominsky and Scholl (2020) by showing that the visual system *is* sensitive to a causal event category that is distinct from launching. Whereas past results are compatible with the idea that "causal perception" is really just "launching(-like) perception" and everything else is the product of abstract (nonperceptual) inference, the retinotopically specific adaptation to entraining, and in particular *causal* entraining (and not the one-object event), shows that there are truly multiple causal *perceptions*.

Of course, this interpretation depends on these effects being genuine visual adaptation effects, rather than something like a response bias. Such concerns are not trivial, as there are some response effects that have been reported for judgments of launching events. One study looking at causal judgments of launching events with variable delays at the moment of contact found a strong sequence effect, such that the response curve for an individual trial was consistently affected by whether the participant responded "causal" on the previous trial (Deodato & Melcher, 2022). However, a response bias account does not explain the retinotopic specificity of the visual adaptation effects reported here and in previous work. A response bias account would basically say that after seeing a bunch of events that clearly belong to one category, participants' default response would be to report seeing the other category. However, there is no reason that such a shift would be stronger at the retinotopic adaptation location. Indeed, if anything one might expect it to be stronger at the spatiotipic adaptation location, if the processing behind the effect was based in a cognitive process that is unconcerned with where the information arrives on the retina.

5.1. *What is being adapted?*

Because all four adaptation conditions in Experiment 1 showed an adaptation transfer effect, we still do not have a precise characterization of *what*, exactly, is being adapted in the launch/pass contrast. We have ruled out the broadest possible interpretation ("causality" as an abstract concept) as well as several versions of the narrowest interpretation (a single low-level visual feature that may not be specific to causal events, or conversely, each event with a distinct verbal description as its own category in vision). At the same time, we are not yet able to present a complete positive account of the minimal set of necessary and sufficient features required to get this adaptation effect for launch/pass displays. (Note that we discuss a "minimal set" because the diversity of events that produce this effect strongly suggests that no single isolated feature is likely to drive these effects, or they are driven by a "single" feature that is itself a composite like "role assignment.")

One direction for future work would be to more exhaustively test combinations of features of the launching event to isolate the minimal set required to get a launch/pass adaptation effect. The most critical question is whether the minimal set of features can be implemented in an event that adults would not describe as "causal." However, given the results to date, we believe it is unlikely that one could get this adaptation effect from a noncausal event, for two reasons. First, something like "pseudo-elastic collision causality" or "launching-like causality" seems to us to be the most parsimonious description for the pattern of events that generate this adaptation effect in past work and Experiment 1. Second, it seems entirely plausible that part of the basis of the strong causal phenomenology produced by these events (see also Supplementary Experiment A1) is the visual processing that is subject to adaptation. Put differently, the impression of "causality" that people mention when describing these events is partially (but not wholly) dependent on the detection of a specific set of features which is subject to adaptation, similar to how the impression of "a face" is partially (but not wholly) dependent on a set of features in particular configuration (Block, 2023; Hochmann & Papeo, 2021). If so, any event that has the relevant set of features (and will, therefore, generate adaptation transfer) will "look" causal and be described as such, even though there are also events that might not have this set of features but are still described as causal (e.g., entraining).[2]

A further wrinkle on the question of what exactly is being adapted in launch/pass displays has to do with the retinotopic nature of the adaptation effects. First, there is some disagreement about whether these effects are genuinely retinotopic (Arnold, Petrie, Gallagher, & Yarrow, 2015). Indeed, both in the present experiments and in some past work (Kominsky and Scholl, 2020), there is clear evidence of some spatiotopic or nonspecific adaptation (i.e., the blue lines in Fig. 3 are all above 0), but it is consistently weaker than the adaptation effect at the retionotpic adaptation location.[3] It is possible in principle that there are some features of launching which are processed in a retinotopic frame of reference, while others are processed in a nonspecific frame of reference. Notably, Ohl & Rolfs (2025) did not test for retinotopic specificity, so it is unknown whether the direction-specificity they reported is across the whole visual field or just for the retinotopic area of adaptation, and conversely, whether the insensitivity to object identity and color is retinotopically specific or not. As such, one could perhaps find an adaptation stimulus that removes some features of launching and shows only nonspecific adaptation, but adaptation nonetheless (e.g., the "gap launch" in Gallagher & Arnold, 2018).

If there are some features which are processed retinotopically and others which are not, it could provide more granular insights into the actual computations the visual system is performing to detect launching events by teasing apart which features seem to be processed together and which ones are processed separately. However, in cases discussed previously that do *not* show retinotopically specific adaptation (e.g., adaptation to "slip" events, or adaptation to entraining events with the launch/pass contrast), there is no adaptation at *either* location, rather than there being equal adaptation at both (Kominsky and Scholl, 2020; Rolfs et al., 2013). As such, it is possible that the particular combination of features that drives this adaptation effect is processed primarily in a retinotopically specific frame of reference. Alternatively, as discussed above, nonspecific adaptation effects may suggest a response bias rather than a genuine perceptual effect, whereas features subject to retinotopically specific adaptation would indicate a perceptual effect above and beyond any response bias. In either case, an important consideration for future work is that we should test for retinotopic specificity whenever we are trying to isolate a particular feature of causal events.

There are similar open questions about the nature of the launch/push distinction, and what exact features of "entraining" drive the adaptation effect in Experiments 2 and 3. The one-object condition rules out "continuous movement of A" as the feature driving the effect. However, the one-object event differs from entraining in several ways, most obviously the presence, contact with, and motion onset of object B. There are events which could be used as adaptation streams to test each whether each of these features drives the entraining adaptation effect. However, given that the nature of the launch/push distinction is based on *when A stops moving*, we find these alternative features to be far less likely candidates to drive the adaptation effect than the motion of A alone. Given that the motion of A alone does not yield the adaptation effect, the entraining event as a whole seems to us to be the most parsimonious explanation. However, future work that more exhaustively tested adaptation streams that vary the different features of entraining events the same way Experiment 1 systematically varied different features of launching events could provide a more precise characterization of the set of features that the visual system uses to identify this category of "entraining."

Another intriguing possibility is that the *launching* adaptation effect found in Experiment 2 may rely on different features than the launch/pass adaptation effect in Experiment 1 and previous work. Consider the following analogy: If a participant is adapted to a Gabor patch made up of different shades of red and animated to give the impression of unambiguous downward motion, we would expect to find either a color aftereffect or a motion aftereffect depending on what kind of test stimulus we used, but the adaptation is not to "red downward motion" as an atomic whole. Rather, the visual system adapts to color, and it adapts to motion direction, and this particular stimulus happens to have both. Similarly, if launching adaptation is adaptation to a combination of features, then the launch/pass and launch/push adaptation effects could be driven by different subsets of those features.

We cannot say for certain whether the features of launching that generate adaptation effects in launch/pass contrasts are the same features that generate adaptation in launch/push contrasts. We did not test adaptation transfer from the events that show adaptation transfer for the launch/pass distinction (tool effect, bursting, and state change; Experiment 1), nor did we test the parametric variations that have demonstrated adaptation transfer to launch/pass displays in past work (absolute and relative speed, color; Kominsky and Scholl, 2020; Ohl & Rolfs, 2025). Notably, for entraining, whatever features produce the adaptation effect for the launch/push contrast *do not* produce an adaptation effect for the launch/pass contrast. Thus, it is very possible that the adaptation to "launching" for launch/pass displays and the adaptation to "launching" for launch/push displays are actually adaptation to different sets of features that are both present in launching. This too would be a worthwhile question for future investigation, as it could provide insight into the separate computations the visual system must perform over different feature dimensions to produce the impression of "cause and effect."

Finally, there are some open questions which may not be possible to address with the current approach because of the limitations of the adaptation paradigm itself. In particular, the adaptation paradigm requires a test event that is ambiguous based on a continuous feature dimension. For the launch/pass display, the difference is about whether the event is seen as causal. For launch/push, it is which of two causal events it is seen as. In both cases, the perceptual experience can be manipulated along a continuous feature dimension. However, if one wanted to determine whether the adaptation transfer effects found in Experiment 1 are symmetrical, for example, whether adapting to launching affected the perception of bursting, one would first need to come up with an ambiguous event that can look like bursting or something else (ideally noncausal), and can be similarly manipulated along a single feature dimension. We have no idea as to what such an event would look like in the case of tool effect, bursting, or state change events; the overlap manipulation of launch/pass would not work, nor would the co-movement of launch/push, and temporal delay introduces confounds with exposure time. Unless or until someone devises a suitable ambiguous versions of these events, the adaptation paradigm cannot be used to look for aftereffects *on* tool effect, bursting, or state change events, only the effects they produce on other events.

5.2. What does this mean for causal cognition?

Thus far, the focus has been on the perceptual nature of these effects, but what makes causal perceptions so interesting as a phenomenon because they suggest the presence of an abstract relation in relatively low-level automatic visual processing (Hafri & Firestone, 2021). However, it is very obviously not the case that our *entire* understanding of causality is visual perception alone. To start with, it is clear that we can make distinctions that are much finer than the categories demonstrated through adaptation; just consider the fact that we have different terms for each of the adaptation events used in Experiment 1. The evidence reported here cannot tell us whether these further distinctions are made in perception or not, only that they are not subject to retinotopically specific visual adaptation on the launch/pass contrast dimension. However, there is ample other work that suggests that our causal *representations* are not entirely based in perception, and in fact may diverge from perception very early on.

While there is some argument that launching is a special category to the human mind from birth (Mascalzoni, Regolin, Vallortigara, & Simion, 2013), just because it is a special event category does not mean it is intrinsically understood as *causal*. To say that an event is understood as causal implies two things: First, that it is distinguished in some way from noncausal events, and second, that its representation in the mind has some kind of causal *content* that supports inference (Kominsky & Carey, 2024). By those criteria, infants do not understand launching as a genuinely *causal* event until they are at least 6 months old, when they show sensitivity to reversals of causal roles (Bélanger & Desrochers, 2001; Desrochers, 1999; Leslie & Keeble, 1987). However, there is already evidence that causal roles *cannot* be tracked through the same perceptual mechanisms that are subject to visual adaptation, since reversing the colors (and thereby the identities and causal roles) of the objects between adaptation and test does not disrupt the adaptation effect (Ohl & Rolfs, 2025).

This is not the only divergence with infant work. While we found evidence that "state change" events adapt launch/pass displays similarly to launching events, infants make different inferences about these two types of events. In particular, in direct comparisons between the two, infants expect that launching should involve contact regardless of whether the causal agent is an inanimate artifact or an animate agent, but their contact expectations for state change events are specific to animate agents (Adibpour & Hochmann, 2023; Muentener & Carey, 2010). Notably, these findings have to do with the causal inferences that infants make from these events. We can, therefore, propose the following account of these divergences: Causal perceptions might solve the first half of the problem of causal understanding by distinguishing causal from noncausal events (and distinguishing some causal events from each other), but the perceptual processes we have identified here may not (on their own) provide the causal content that supports more domain-general inference processes (and what *does* provide that content is not something we have space to discuss here). To be very precise, what we are proposing is that humans perceive "a set of features inseparable from causality but that has no inference-supporting causal content." For brevity, even if this supposition should prove to be true, we recommend the field continue to use "causal perceptions" as a shorthand.

5.3. Conclusion

Breaking monolithic "causal perception" into a set of feature dimensions that are subject to adaptation does not contradict the idea that perception encodes higher-level relations (Hafri & Firestone, 2021; Papeo, 2020; Scholl & Tremoulet, 2000), but nor does it support the notion that perceptual representations of relations support abstract propositional content (Hochmann & Papeo, 2021). That is not to say that we regard our results as any kind of deflationary account of what the visual system is doing, or counter to the idea that the visual system is tuned to certain types of relations in the environment. On the contrary, the sets of features adapted in these studies are (thus far) specific to causal relations between objects, and seem likely to be complex combinations of features rather than any one easily-isolated property of these events that is readily found in a noncausal event. The current state of the evidence indicates that these are, in at least one sense, genuine causal *perceptions*, that is, specialized visual processing routines that are (thus far) exclusively sensitive to causal events. However, the perceptual systems that we can identify through adaptation may only pick out the events that have these combinations of features, and leave it to more generalized downstream processing to provide representations of these events with richer conceptual or propositional content (e.g., agent and patient roles). As discussed above, the identification and comprehension of causal interactions in our environment is clearly a rich and multilayered process, but the distinct causal perceptions we can uncover using visual adaptation reveal some early "joints" that the mind uses to carve up our experience of the dynamic world that surrounds us.

Acknowledgments

This research was supported by an internal CEU Research Support Scheme grant to JFK.

Open access funding provided by Central European University Private University - CEU GmbH/KEMÖ.

Notes

1 There is one possible piece of counterevidence to this, which is that Gallagher and Arnold (2018) reported getting an adaptation effect with an event where A stopped one object radius short of B, that is, did not make contact. However, it is difficult to draw firm conclusions from this result as (1) the adaptation effect observed in this case was not retinotopically specific (see also General Discussion) and (2) there is some evidence that a small spatial offset at high speeds still looks like causal launching (Michotte, 1963, experiments 31–32, pp. 99–100).

2 Notably, there are some judgments that people make about these events which are likely *not* based on perception. For example, judgments of "force" based on the dynamics of bursting/shattering events (Hubbard & Ruppel, 2013) can be replicated with written descriptions of the stimuli alone (van Buren, & Scholl, 2025). However, our results argue that shattering may still be identified as belonging to the category of *causal* events in early perceptual processing.

3 One could also argue that the effects reported here and in Kominsky and Scholl (2020) are hemifield-specific rather than retinotopically specific, since the test locations were always left and right of the fixation dot (this does not apply to Rolfs et al., 2013). It may be difficult to fully tease these possibilities apart, because the receptive field size of the brain areas that are likely to be involved (e.g., motion areas like V5/MT) is quite large (Kolster et al., 2010). However, for our current purposes, the distinction is not of paramount importance, because retinotopic specificity is only used as evidence of informationally encapsulated perceptual processing, and a hemifield-specific effect would support the same conclusions. This issue will be more relevant for future work that wishes to examine the neural and computational underpinnings of these effects in more detail.

References

Addams, R. (1834). An account of a peculiar optical phenomenon seen after having looked at a moving body. *Cognitive Development*, *5*, 373–374.

Adibpour, P., & Hochmann, J.-R. (2023). Infants' understanding of the causal power of agents and tools. *Proceedings of the National Academy of Sciences*, *120*(50), e2309669120. https://doi.org/10.1073/pnas.2309669120

Arnold, D. H., Petrie, K., Gallagher, R., & Yarrow, K. (2015). An object-centered aftereffect of a latent material property: A squishiness visual aftereffect, not causality adaptation. *Journal of Vision*, *15:4*, 1–9. https://doi.org/10.1167/15.9.4

Benton, D. T. (2024). An associative-learning account of how infants learn about causal action in animates and inanimates: A critical reexamination of four classic studies. *Journal of Experimental Psychology: General*. https://doi.org/10.1037/xge0001656

Bélanger, N. D., & Desrochers, S. (2001). Can 6-month-old infants process causality in different types of causal events? *British Journal of Developmental Psychology*, *19*(1), 11–21. https://doi.org/10.1348/026151001165930

Block, N. (2023). *The border between seeing and thinking*. Oxford University Press.

Bock, O., Baetge, I., & Nicklisch, A. (2014). hroot: Hamburg registration and organization online tool. *European Economic Review*, *71*, 117–120. https://doi.org/10.1016/j.euroecorev.2014.07.003

Choi, H., & Scholl, B. (2006). Perceiving causality after the fact: Postdiction in the temporal dynamics of causal perception. *Perception*, *35*(3), 385–399. https://doi.org/10.1068/p5462

Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of the causality of a simple event. *Infant Behavior and Development*, *21*(4), 713–731. https://doi.org/10.1016/S0163-6383(98)90040-6

Cohen, L. B., Rundell, L. J., Spellman, B. A., & Cashon, C. H. (1999). Infants' perception of causal chains. *Psychological Science*, *10*(5), 412–418. https://doi.org/10.1111/1467-9280.00178

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

Desrochers, S. (1999). Infants' processing of causal and noncausal events at 3.5 months of age. *Journal of Genetic Psychology*, *160*(3), 294–302. https://doi.org/10.1080/00221329909595400

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, *43*(9), 1035–1045. https://doi.org/10.1016/s0042-6989(03)00084-1

Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.

Gallagher, R. M., & Arnold, D. H. (2018). Comparing the aftereffects of motion and causality across visual coordinates. *BioRxiv*, 499764. https://doi.org/10.1101/499764

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

Hochmann, J.-R., & Papeo, L. (2021). How can it be both abstract and perceptual? Comment on Hafri, A., & Firestone, C. (2021), the perception of relations, trends in cognitive sciences. *PsyArXiv*. https://osf.io/preprints/psyarxiv/hm49p/download

Hubbard, T. L. (2013). Phenomenal causality I: Varieties and variables. *Axiomathes*, *23*(1), 1–42. https://doi.org/10.1007/s10516-012-9198-8

Hubbard, T. L., & Favretto, A. (2003). Naïve impetus and Michotte's "tool effect": Evidence from representational momentum. *Psychological Research*, *67*(2), 134–152. https://doi.org/10.1007/s00426-002-0122-5

Hubbard, T. L., & Ruppel, S. E. (2013). Ratings of causality and force in launching and shattering. *Visual Cognition*, *21*, 987–1009. https://doi.org/10.1080/13506285.2013.847883

Karaminis, T., Turi, M., Neil, L., Badcock, N. A., Burr, D., & Pellicano, E. (2015). Atypicalities in perceptual adaptation in autism do not extend to perceptual causality. *PLoS One*, *10*(3), e0120439. https://doi.org/10.1371/journal.pone.0120439

Kolster, H., Peeters, R., & Orban, G. A. (2010). The retinotopic organization of the human middle temporal area MT/V5 and its cortical neighbors. *Journal of Neuroscience*, *30*(29), 9801–9820. https://doi.org/10.1523/JNEUROSCI.2069-10.2010

Kominsky, J. F., & Carey, S. (2024). Infants' representations of Michottean triggering events. *Cognition*, *250*, 105844. https://doi.org/10.1016/j.cognition.2024.105844

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

Kominsky, J. F., Strickland, B., Wertz, A., Elsner, C., Wynn, K., & Keil, F. (2017). Categories and constraints in causal perception. *Psychological Science*, *28*(11), 1649–1662. https://doi.org/10.1177/0956797617719930

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, *25*(3), 265–288. https://doi.org/10.1016/S0010-0277(87)80006-9

Liu, S., Brooks, N. B., & Spelke, E. S. (2019). Origins of the concepts cause, cost, and goal in prereaching infants. *Proceedings of the National Academy of Sciences*, *116*(36), 201904410. https://doi.org/10.1073/pnas.1904410116

Mascalzoni, E., Regolin, L., Vallortigara, G., & Simion, F. (2013). The cradle of causal reasoning: Newborns' preference for physical causality. *Developmental Science*, *16*(3), 327–335. https://doi.org/10.1111/desc.12018

Mather, G., Pavan, A., Campana, G., & Casco, C. (2008). The motion aftereffect reloaded. *Trends in Cognitive Sciences*, *12*(12), 481–487. https://doi.org/10.1016/j.tics.2008.09.002

Mayrhofer, R., & Waldmann, M. R. (2014). Indicators of causal agency in physical interactions: The role of the prior context. *Cognition*, *132*(3), 485–490. https://doi.org/10.1016/j.cognition.2014.05.013

Michotte, A. (1963). *The perception of causality*. New York: Basic Books.

Michotte, A. (1991). The perception of the "tool effect". In G. Thinès, A. Costall, & G. Butterworth (Eds.), *Michotte's experimental phenomenology of perception* (pp. 87–103). Hillsdale, NJ: Routledge

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

Muentener, P., & Carey, S. (2010). Infants' causal representations of state change events. *Cognitive Psychology*, *61*(2), 63–86. https://doi.org/10.1016/j.cogpsych.2010.02.001

Natsoulas, T. (1961). Principles of momentum and kinetic energy in the perception of causality. *American Journal of Psychology*, *74*(3), 394–402. https://doi.org/10.2307/1419745

Ohl, S., & Rolfs, M. (2025). Visual routines for detecting causal interactions are tuned to motion direction. *eLife*, *13*, RP93454. https://doi.org/10.1101/2023.08.22.554237

Papeo, L. (2020). Twos in human visual perception. *Cortex*, *132*, 473–478. https://doi.org/10.1016/j.cortex.2020.06.005

Papeo, L., Vettori, S., Serraille, E., Odin, C., Rostami, F., & Hochmann, J.-R. (2024). Abstract thematic roles in infants' representation of social events. *Current Biology*, *34*(18), 4294–4300.e4. https://doi.org/10.1016/j.cub.2024.07.081

Peirce, J., Gray, J., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. (2019). Psychopy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*, 195–203. https://doi.org/10.3758/s13428-018-01193-y

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

Saxe, R., & Carey, S. (2006). The perception of causality in infancy. *Acta Psychologica*, *123*(1–2), 144–165. https://doi.org/10.1016/j.actpsy.2006.05.005

Scholl, B., & Nakayama, K. (2002). Causal capture: Contextual effects on the perception of collision events. *Psychological Science*, *13*(6), 493–498. https://doi.org/10.1111/1467-9280.00487

Scholl, B., & Tremoulet, P. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, *4*(8), 299–309. https://doi.org/10.1016/S1364-6613(00)01506-0

Ullman, S., Harari, D., & Dorfman, N. (2012). From simple innate biases to complex visual concepts. *Proceedings of the National Academy of Sciences*, *109*(44), 18215–18220. https://doi.org/10.1073/pnas.1207690109

van Buren, B. F., & Scholl, B. J. (2025). The blindfold test: Helping to decide whether an effect reflects visual processing or higher-level judgment. *Attention, Perception, & Psychophysics*, *87*(2), 445–457. https://doi.org/10.3758/s13414-024-02939-x

Vettori, S., Odin, C., Hochmann, J.-R., & Papeo, L. (2025). A perceptual cue-based mechanism for automatic assignment of thematic agent and patient roles. *Journal of Experimental Psychology. General*, *154*(3), 787–798. https://doi.org/10.1037/xge0001657

Webster, M. A. (2016). Visual adaptation. *Annual Review of Vision Science*, *1*, 547–567. https://doi.org/10.1146/annurev-vision-082114-035509

White, P. A. (2009). Perception of forces exerted by objects in collision events. *Psychological Review*, *116*(3), 580–601. https://doi.org/10.1037/a0016337

White, P. A., & Milne, A. (1999). Impressions of enforced disintegration and bursting in the visual perception of collision events. *Journal of Experimental Psychology: General*, *128*(4), 499.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Supporting Information