Attention, Perception, & Psychophysics (2024) 86:559–566
https://doi.org/10.3758/s13414-023-02837-8

Spatiotemporal jump detection during continuous film viewing: Insights from a flicker paradigm

Aditya Upadhyayula[1] · **John M. Henderson**[2,3]

Accepted: 19 December 2023 / Published online: 3 January 2024

Abstract
We investigated how sensitive visual processing is to spatiotemporal disruptions in ongoing visual events. Prior work has demonstrated that participants often miss spatiotemporal disruptions in videos presented in the form of scene edits or disruptions during saccades. Here, we asked whether this phenomenon generalizes to spatiotemporal disruptions that are not tied to saccades. In two flicker paradigm experiments, participants were instructed to identify spatiotemporal disruptions created when videos either jumped forward or backward in time. Participants often missed the jumps, and forward jumps were reported less frequently compared with backward jumps, demonstrating that a flicker paradigm produces effects similar to a saccade contingent disruption paradigm. These results suggest that difficulty detecting spatiotemporal disruptions is a general phenomenon that extends beyond trans-saccadic events.

Keywords Visual cognition · Spatiotemporal disruptions · Change blindness · Flicker paradigm · Film comprehension

Introduction

Our visual system is constrained by bottlenecks at various stages of processing. As a result, we cannot process all of the information that enters our retina. A large body of work has investigated bottlenecks at various stages of visual processing, including but not limited to retinal organization (Carrasco et al., 2005; Kolb, 2011; Rosenholtz, 2016; Upadhyayula et al., 2023), attention (Drew et al., 2013; Henderson & Hollingworth, 2003; Hollingworth & Henderson, 2000; Posner, 1980; Scholl, 2000; Smith et al., 2012; Wolfe & Gray, 2007), eye movements (Fehd & Seiffert, 2008; Henderson & Hollingworth, 1998; Rayner, 1975, 1998; Upadhyayula & Flombaum, 2020; Zelinsky, 2001; Zelinsky & Neider, 2008), and crowding (Levi, 2008; Whitney & Levi, 2011). Yet, our experience of the visual world seems rich and continuous. Here, we investigated whether the visual system faces similar bottlenecks in its sensitivity to spatiotemporal disruptions in naturalistic dynamic visual events.

Surprisingly, little research has investigated sensitivity to visual spatiotemporal continuity (Magliano et al., 2001; Magliano & Zacks, 2011; Smith & Henderson, 2008; Upadhyayula & Henderson, 2023). In one study, Smith and Henderson (2008) demonstrated that participants frequently failed to detect film edits, a phenomenon they called *edit blindness*. The results demonstrated that edit blindness is greater for edits during an ongoing action compared with edits between actions. Magliano and colleagues (Magliano et al., 2001) demonstrated that situational continuity—such as changes in spatial movement, location and time—were important factors in constructing events during film comprehension. In the same vein, Magliano and Zacks (2011) showed that discontinuity in movies at event boundaries involving situational discontinuities produced distinct patterns in the primary visual cortex that were significantly different from other types of edits. Based on this result, they concluded that specialized mechanisms in higher order perceptual processing could be involved in maintaining an active representation of film with discontinuities. These results collectively suggest that sensitivity to spatiotemporal continuity in visual processing is limited. However, this research primarily focused on

✉ Aditya Upadhyayula
aditya.usa8@gmail.com

1 Department of Psychological & Brain Sciences, Washington University in St. Louis, CB 1125, One Brookings Drive, St. Louis, MO 63130-4899, USA

2 Center for Mind and Brain, University of California, Davis, USA

3 Department of Psychology, University of California, Davis, USA

scene transitions at event boundaries or at edits that were particularly tailored to the narrative.

More recently, Upadhyayula and Henderson (2023) reported that blindness to disruptions in video extends beyond scene edits. In this study, participants' eye movements were recorded as they watched movie clips that did not include scene edits. Occasionally during saccades, the videos either jumped forward—that is, the video moved forward in time more than the passage of time warranted, or backward (i.e., the video moved backward in time so that a portion of the video was seen again). Their data showed that participants often missed the jumps, even when the jumps were as large as 2,000 ms. Furthermore, participants were less sensitive to detecting forward jumps compared with backward jumps. Based on the asymmetry in jump direction, with lower detection rates for large forward jumps, the study suggested that knowledge about an unfolding event could potentially make spatiotemporal jump detection more difficult.

Although the sensitivity to trans-saccadic spatiotemporal disruptions observed by Upadhyayula and Henderson (2023) was taken to index a general property of spatiotemporal processing, it is possible that the results were related to the eye movements themselves. For example, the disruptions were tied to each participant's specific eye movements, which were uncontrolled by the experimenters and therefore not randomly distributed in time or space. Participants' eye movements are influenced by the visual content of the viewed image and how the viewer interprets that content, and it is therefore possible that some feature of the videos could have influenced both saccadic movements and change detection. To decouple any potential saccade-specific effects from general spatiotemporal disruptions, here we used a flicker paradigm in which the presented disruptions took place at experimenter-determined times during flickers rather than participant-determined times during saccades.

Flicker paradigms have been widely used in the literature to study change blindness in static images (Henderson et al., 2008; Henderson & Hollingworth, 1999; Hollingworth & Henderson, 2000; Rensink et al., 1997; Scholl, 2000; Simons, 2000; Simons & Levin, 1997; Simons & Rensink, 2005). In the flicker paradigm, an original image and a changed image are presented in rapid alternation in time separated by a brief blank image. Observers are instructed to respond as soon as they detect the change. The underlying idea is that the blank image briefly disrupts continuous visual processing. Furthermore, flickers are thought to impair the local motion signals that accompany the change (Simons & Rensink, 2005). Results from the flicker paradigm have been comparable (though not identical) to those based on saccadic eye movements demonstrating change blindness (Henderson et al., 2008).

The present study used a flicker paradigm to disrupt the spatiotemporal continuity of dynamic image (video) in the same way as saccadic eye movements did in Upadhyayula and Henderson (2023). To our knowledge, this is the first study to use the flicker paradigm in video to study how the visual system processes spatiotemporal continuity. Unlike the traditional flicker paradigm where the original and the modified images are rapidly alternated, the videos in our experiments either jumped forward or backward once during a predetermined flicker. To maintain continuity in video comprehension, our visual system must overcome these disruptions by bridging the gaps during disruptions caused by the flickers. If the results observed in our previous saccade-contingent jump study generalize beyond saccades, we expect insensitivity to spatiotemporal discontinuity induced by the jumps. Furthermore, we hypothesize that our visual system relies on the knowledge of unfolding information to bridge the gaps in spatiotemporal continuity. Therefore, any disruptions resulting in a change in the video in the direction of unfolding knowledge should be less likely to be noticed. Consequently, we should expect less sensitivity to forward jumps compared with backward jumps. Put differently, the observed detection rates should be lower for forward jumps compared with backward jumps.

Experiment 1

Methods

Participants

Seventy-six undergraduates participated in this study online for course credit. We sought to test a minimum of 50 participants. Upon completion of the study, participants completed a Likert scale questionnaire, ranging between 1 and 5 (5 being the highest), regarding internet connectivity during the study, how well they understood the instructions, and study difficulty. Ten participants were eliminated based on response ratings less than or equal to 1, thus leaving 66 participants in total. The protocols for the reported experiments were approved by the Institutional Review Board (IRB) of UC Davis. Participant eye movements were neither enforced nor monitored during the study.

Stimuli and apparatus

We used the stimuli from Upadhyayula and Henderson (2023). The stimuli were 36 one-minute video clips without audio from the film *1917*. This film was chosen because it does not have any perceivable scene edits. All video clips contained camera panning/zooming. The video clips were generated using ffmpeg to separate audio and video streams from the original movie clips, and the audio stream was discarded. The extracted video streams were further compressed and resized to 463 × 240 pixels, with a bit rate of 24 frames per second to facilitate seamless video playback at lower internet speeds. Finally, two additional versions of each video clip were generated with the video clips starting either 500 ms or 1,000 ms later than the start of the original clip. This way, each video clip had two delayed versions of the same clip.

Stimuli were presented using JsPsych 7.0.0, a JavaScript library for creating online behavioral experiments (de Leeuw, 2015). A custom JsPsych plugin was programmed to create the flicker paradigm (see below for a demo). Stimuli were presented on web browsers at a resolution of 1,000 × 518 pixels, thus preserving the original aspect ratio. Participants performed the task on their personal computing machines.

Design and procedure

A demo of the experiment can be viewed here (https://adibuoy23.github.io/temporal_change_blindness_longer_shifts/video_cb_flicker_paradigm_longer_shifts.html). Each trial began with a central fixation point on the screen on a white background. All three versions of the video clip (original and the delayed versions) started playing simultaneously 1,000 ms after the central fixation. However, only one video was visible at any given time. Furthermore, a white mask equal to the size of the video was periodically displayed on top of the video every 2,000 ms for a duration of 150 msec during which the video was rendered invisible. This created an effect of the flicker during video playback. Occasionally during these flickers, the visibility of the original and the delayed video clips was toggled one at a time to create a jump. On average, the jumps were between 4 and 6 seconds apart. In this way participants could not predict which flickers would contain jumps. The specific flickers when the videos jumped were randomly selected prior to the experiment such that all the participants experienced the jumps at the same time points. Overall, these manipulations ensured a seamless video that jumped by either 500 or 1,000 ms during occasional flickers.

Each change was randomly selected with the constraint that changes always occurred in pairs: A change in one direction was always followed by a return change in the opposite direction and of an equal magnitude during a subsequent flicker. This was done to balance the number of forward and backward changes and magnitudes per trial. Each trial contained five changes: one backward 1000 ms (where participants had to re-watch the last second in the clip), one backward of 500 ms, one forward of 1,000 ms (where participants skipped ahead into the video by 1 second), one forward 500 ms. A 0-ms change was also used occasionally during the flickers where the program did not toggle between the videos. Any responses made during the 0 ms and the nonchange flickers were used to analyze the false alarms while detecting jumps in the videos.

Participants were instructed to press the SPACEBAR whenever they noticed any jumps in the video. Responses were recorded as correct if they responded within 1,750 ms of a change during the flicker. Responses beyond this cutoff were not recorded and thus excluded from analysis. Each trial lasted about one minute. Participants performed one practice trial with feedback—where they were told when they detected and missed the jumps. They also had an opportunity to redo the trial to get a better understanding of the paradigm. Feedback was only provided during the practice and was disabled for the rest of the experiment. Figure 1 displays the schematic of the experiment.

Data analysis and results

The materials used for data analysis are available via the OSF repository (https://osf.io/296jh/). We analyzed participant keypresses in the primary analysis. Participant key presses were coded as 1 if they pressed the button within 1.75 seconds of the display change, 0 if they missed it. Practice trials were excluded from analysis, leaving 35 experimental trials in total for each participant. Analysis was performed in the R programming environment (Version 4.1.1; R Core Team, 2019). Generalized Linear mixed-effects (GLME) models from the R package "lmerTest" (Version 1.1-28; Bates et al., 2015), were fit to the data given the categorical nature of the keypresses. Change direction (forward vs. backward) and magnitude (500 vs. 1,000 ms) were the predictor variables. Trials and participants were treated as random effects with direction and magnitude accounted for by random slopes.

On average, participants detected the video jumps about 50% of the time ($M = 0.49$, $SD = 0.5$). Forward jumps were detected about 44% of the time ($M = 0.44$, $SD = 0.49$) which is less than the 54% detection rate in the backward jumps ($M = 0.54$, $SD = 0.49$). Furthermore, 500-ms jumps were detected about 44% of the time ($M = 0.44$, $SD = 0.4$) which is less than the 55% detection rate in the 1,000 ms jumps ($M = 0.55$, $SD = 0.49$). The raw data are shown in Fig. 2. GLME analysis revealed a significant effect of change direction (odds ratio = 0.59, $p < 0.001$), significant effect of change magnitude (odds ratio = 1.78, $p < 0.001$). The interaction between change direction and magnitude was not significant (odds ratio = −1.23, $p = 0.26$), see Fig. 3 and Table 1. We used the flicker instances where there was no change and where participants did not press the button as the true negative reports (correct rejections). A signal detection analysis revealed a $d' = 0.43$; the false-alarm rate was 20.9%, which is about 42% of the average detection rate.

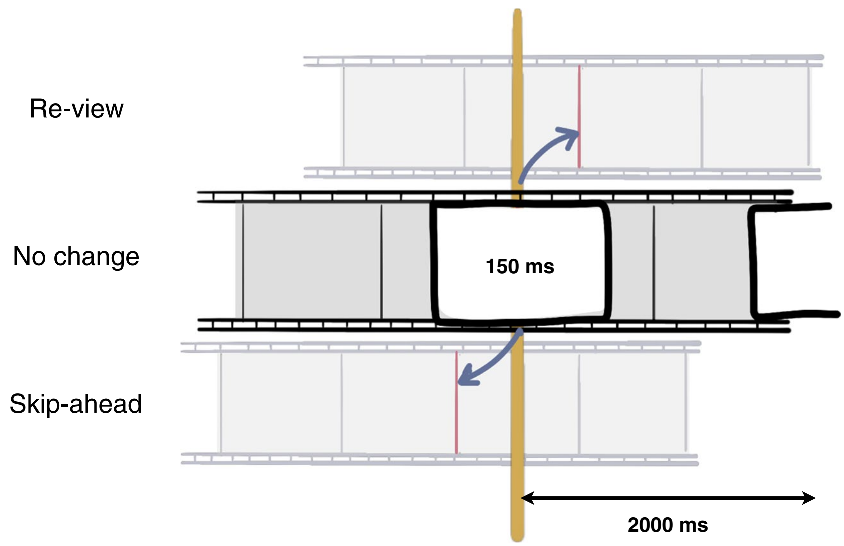

Fig. 1 Schematic of a flicker contingent temporal disruption paradigm. The screen flickered every 2 seconds for a duration of 150 ms. During occasional flickers, the videos either jumped forward in time, or backward

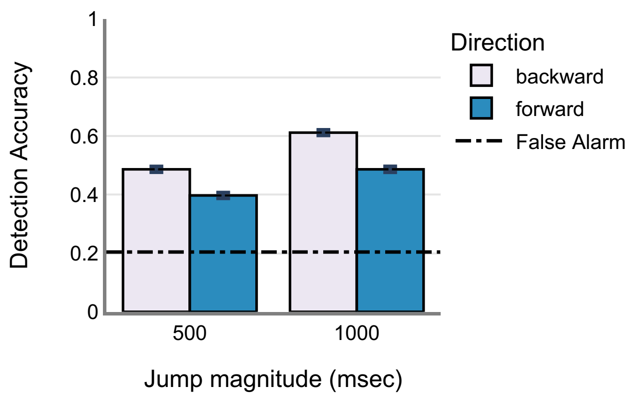

Fig. 2 Observed behavior data from Experiment 1. Error bars indicate standard error of mean (*SEM*)

In summary, the general insensitivity of viewers to spatiotemporal disruptions observed by Upadhyayula and Henderson (2023) across saccades was also observed in a within-fixation flicker paradigm here.

Experiment 2

To determine whether the sensitivity to spatiotemporal jumps extends beyond 1,000 ms, Experiment 2 investigated the boundary conditions of these effects by increasing the magnitude of jumps to 2,000 ms. Experiment 2 was a conceptual replication of Experiment 1, with the modification that the jump magnitude of the videos increased to 1,000 and 2,000 ms.

Methods

Participants

The methods were the same as Experiment 1 with the following exceptions. Seventy-one undergraduates participated online for course credit. Five participants were further eliminated based on their response ratings for understanding the instructions and internet connectivity less than or equal to 1, thus leaving with 66 participants in total. The magnitude of jumps was either 1,000 or 2,000 ms.

Data analysis and results

The data analysis method for this experiment was the same as that of Experiment 1. A preliminary analysis revealed that the random effects from subjects were minimal for the magnitude condition. To avoid singularity and model overfitting, we therefore analyzed the data change direction and magnitude as the fixed effects; trial and subjects as random effects; change direction as the random slope within subjects; and change magnitude, direction, and their interaction as random slopes within the trials.

On average, participants detected the video jumps about 50% of the time (M = 0.52, SD = 0.49). Forward jumps were detected about 43% of the time which is less frequent compared with the 60% accuracy in the backward jumps (forward: M = 0.43, SD = 0.49; backward: M = 0.60, SD = 0.48). Furthermore, 1,000-ms jumps were detected about 48% of the time, which is less frequent compared with the 55% accuracy in the 2,000-ms jumps (1,000 ms: M = 0.48,

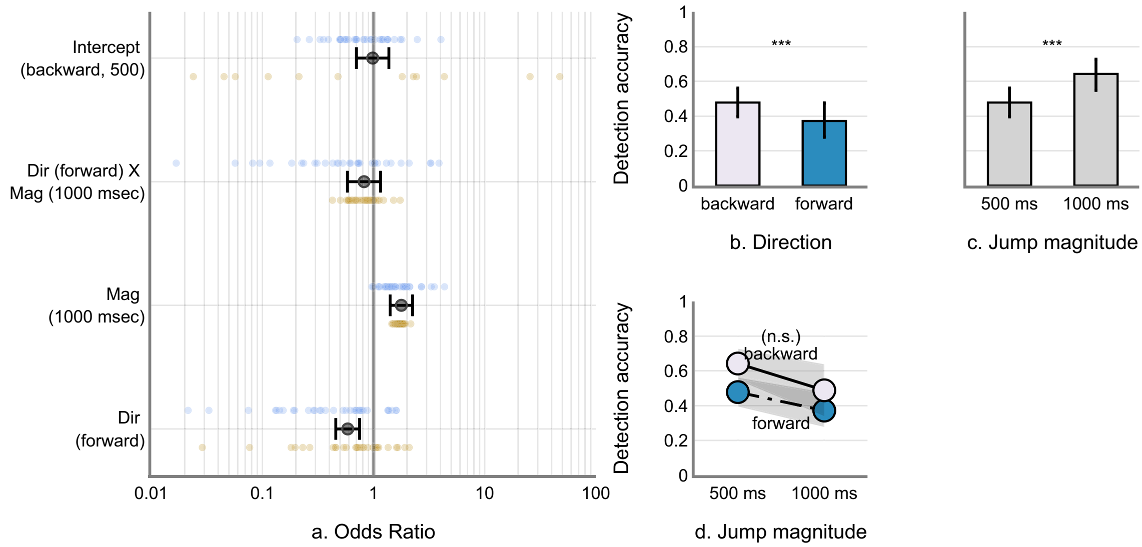

Fig. 3 GLME Model predictions. **a** Model estimates. Blue and yellow dots represent the random effects from the trial and subjects respectively. **b** Marginal estimate of jump direction. **c** Marginal estimate of jump magnitude. d Interaction plot. Error bars and the shaded regions represent 95% CI. (Color figure online)

Table 1 GLME estimates

Predictors	Fixed effects			Random effects, *SD*	
	Odds Ratios	*CI*	*p*	By-trial	By-subject
Intercept (backward, 500 ms)	0.98	[0.70-1.37]	0.916	0.30	1.34
Direction (forward)	0.59	[0.46-0.75]	<0.001	0.25	0.4
Magnitude (1,000 ms)	1.78	[1.41-2.25]	<0.001	0.38	0.06
Direction (forward): Magnitude (1,000 ms)	0.82	[1.58-1.16]	0.26	0.69	0.09

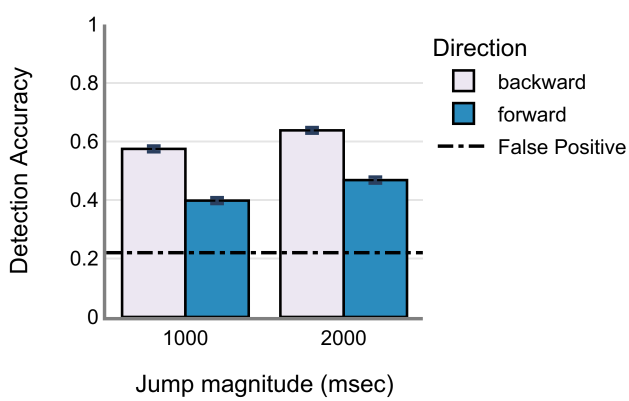

Fig. 4 Observed behavior data in Experiment 2. Error bars indicate standard error of mean (*SEM*)

$SD = 0.49$; 2,000 ms: $M = 0.55$, $SD = 0.49$). The raw data are shown in Fig. 4. GLME analysis revealed a significant effect of change direction (odds ratio = 0.39, $p < 0.001$), significant effect of change Magnitude (odds ratio = 1.45, $p < 0.001$). The interaction between change direction and magnitude was not significant (odds ratio = 1.03, $p = 0.9$); see Fig. 5 and Table 2 . A signal detection analysis revealed a $d' = 0.43$; the false-alarm rate was 21.9%, which is about 43% of the average detection rate.

General discussion

How sensitive is our visual system to spatiotemporal discontinuities while watching videos of naturalistic events? Prior work has characterized the spatiotemporal limitations of film viewing either at scene edits, or during event boundaries (Magliano et al., 2001; Magliano & Zacks, 2011; Smith & Henderson, 2008). The main findings from this literature are that people often miss scene edits (Smith & Henderson, 2008), that situational continuity (e.g., changes in spatial movement, location, and objects) is important for maintaining a coherent representation of the film (Magliano et al., 2001), and that our visual system bridges the gaps in

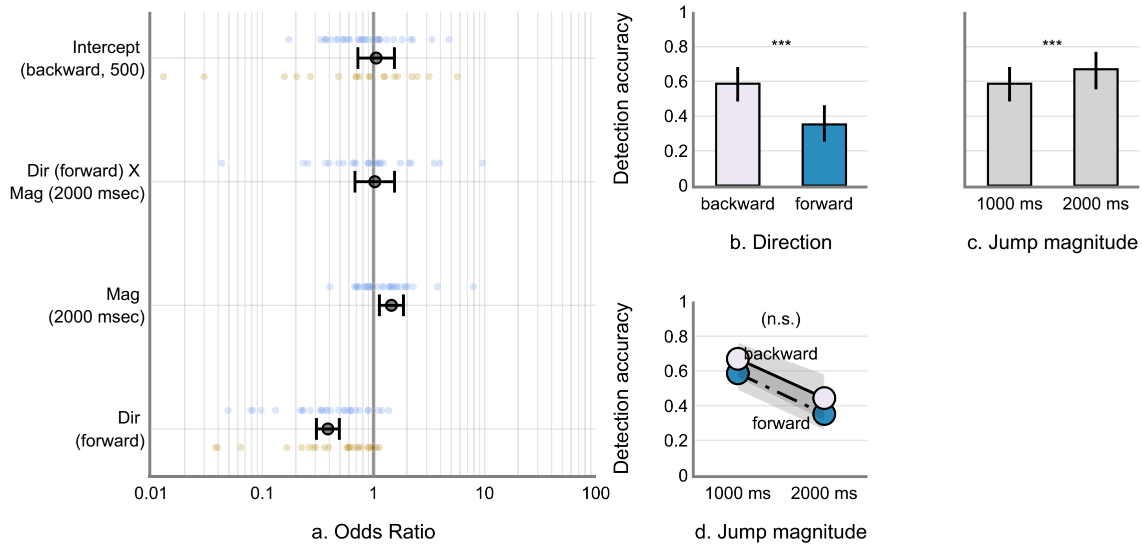

Fig. 5 GLME Model predictions. **a** Model estimates. Blue and yellow dots represent the random effects from the trial and subjects respectively. **b** Marginal estimate of jump direction. **c** Marginal estimate of jump magnitude. **d** Interaction plot. Error bars and the shaded regions represent 95% CI. (Color figure online)

Table 2 GLME estimates

Predictors	Fixed effects			Random effects, *SD*	
	Odds ratios	CI	*p*	By-trial	By-subject
Intercept (backward, 500 ms)	1.06	[0.72, 1.54]	0.777	0.37	1.72
Direction (forward)	0.39	[0.31, 0.49]	<0.001	0.21	0.38
Magnitude (1,000 ms)	1.45	[1.13, 1.86]	<0.001	0.49	-
Direction (forward): Magnitude (1,000 ms)	1.03	[0.68, 1.55]	0.9	1.20	-

film caused during edits and event boundaries by relying on higher level knowledge of the visual scene (Magliano & Zacks, 2011). Whether and how these mechanisms extend beyond scene edits and event boundaries in facilitating the detection of spatiotemporal disruptions is not understood.

Recently, we used eye tracking to demonstrate that insensitivity to spatiotemporal disruptions in videos extends beyond scene edits (Upadhyayula & Henderson, 2023). However, in that study, the disruptions took place during saccadic eye movements. One potential concern is that the time and location of the critical eye movements that generated a jump were uncontrolled by the experimenters. As a result, it is possible that participants had more control over jump detection. Moreover, participant eye movements are rarely random, and could also be influenced by the video content. In the present study, we controlled for this issue using a flicker paradigm in which flicker timing was controlled by the experimenters to investigate spatiotemporal jump detection independently of eye movements. In two experiments, participants performed a flicker version of the spatiotemporal jump detection task and were instructed to report any noticeable jumps in the video clips. The screen flickered every 2 seconds, and occasionally during a flicker the video either jumped forward or backward in time. If the insensitivity to spatiotemporal jumps observed by Upadhyayula and Henderson (2023) is a general phenomenon not specifically tied to eye movements, then we should observe a similar effect here.

Overall, participants detected 50% of the jumps. Furthermore, on average, participants detected the forward jumps 10% less frequently than they did the backward jumps in Experiment 1 (forward = 44%, backward = 54%). Jumps of magnitude 1000 ms were detected 11% more frequently compared with the 500-0ms jumps (500 ms = 44%, 1,000 ms = 55%)—see Fig. 2. This was true even when the magnitude of the jumps was increased in Experiment 2. Overall, forward jumps were detected 17% less frequently than the backward jumps (forward = 43%, backward = 60%). Jumps of magnitude 2,000 ms were detected 7% more frequently than the 1,000-ms jumps (1,000 ms =

48%, 2,000 ms = 55%)—see Fig. 4. These results in conjunction with those observed by Upadhyayula and Henderson (2023) demonstrate that insensitivity to spatiotemporal jumps is not particularly tied to eye movements.

It should be noted that participants in our experiments missed the jumps about 40%–50% of the time on average compared with the 10%–30% miss rate reported in (Upadhyayula & Henderson, 2023). In addition, the measured false alarm rates were on average about 50%–60% of the detection rates in both experiments. These results suggest that jump detection during flickers is harder compared with the jumps during saccades. This is an important finding for at least two reasons: Prior work comparing change detection for saccades contingent changes versus flickers has shown differential sensitivity to changes occurring during saccades compared with flickers (Henderson et al., 2008). Their results showed that flicker changes had higher accuracy compared with the changes happening during saccades. Participants in that study detected changes in luminance and contrast that happened either during saccades or during flickers. Our results show the opposite effect wherein the accuracies for detecting flicker changes were lower compared with the saccade changes from our previous study (Upadhyayula & Henderson, 2023). It is unclear why the sensitivity reversed for flickers in our results. One possible explanation is that the changes presented in Henderson et al. (2008) involved luminance manipulations which were low level in nature, whereas the video jumps in our study could also have involved a change in high-level visual information such as expectations about where the actors in the movie should be after the jump. This could explain the reversed sensitivity for flickers and saccades in the jump detection paradigm. Future work could benefit from further investigation towards this end.

Another related potential reason for the observed difference in detection rates between this study and that reported in Upadhyayula and Henderson (2023) concerns the experiment design. Participants in this study watched videos as the screen flickered every 2 seconds. This type of disruption is forcibly induced compared with the naturally occurring disruptions during saccades. A consequence of such forced disruption in videos is that it could be harder to detect the jumps since flickers are designed to minimize any transient changes, and accordingly relying on low-level visual information for jump detection becomes harder. Indeed, low-level visual properties such as optic flow have been shown to predict the detection of spatiotemporal jumps during saccades (Upadhyayula & Henderson, 2023). It is therefore possible that participants were relying on high-level information to detect spatiotemporal jumps as a result. This could also explain the lower detection rates in this study compared with Upadhyayula and Henderson (2023). Future work could benefit from systematically dissociating the influence of low-level versus high-level visual information in detecting spatiotemporal jumps.

Why were the forward jumps detected less frequently compared with the backward jumps? One possibility is that our visual system makes assumptions about how the future unfolds, and as a result, any expected changes that are in line with such knowledge and expectations are less likely to be noticed. Consistent with this hypothesis, forward jumps were less frequently detected compared with the backward jumps in both the experiments even as the jump magnitude varied from 500 to 2,000 ms—see Figs. 2 and 4. Alternatively, it is possible that backward jumps are easier to detect because the content in the backward jumps has already been seen before. On this hypothesis, rather than postulating that forward jumps are more difficult due to expectations, the focus is on the idea that backward jumps are easier due to memory for the viewed content. However, note that this account would not be able to explain why the detection accuracy increased in the forward jumps as the jump magnitude increased from 500 ms to 2,000 ms—see Figs. 3b and 5b. It is possible that the detection of spatiotemporal jumps is facilitated by a combination of both memory and expectations. Future work could benefit from further investigation towards this end.

Acknowledgments This research was funded by BCS2019445 from the US National Science Foundation. The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The authors declare no competing financial interests.

Data Availability All the materials used for the data analysis are available via the OSF repository (https://osf.io/296jh/).

References

Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). Parsimonious mixed models. *arXiv Preprint arXiv:1506.04967*.

Carrasco, M., Giordano, A. M., & McElree, B. (2005). Temporal dynamics of covert attention. *Journal of Vision, 5*(8), 232–232.

de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods, 47*(1), 1–12. https://doi.org/10.3758/s13428-014-0458-y

Drew, T., Võ, M.L.-H., & Wolfe, J. M. (2013). The invisible gorilla strikes again: Sustained inattentional blindness in expert observers. *Psychological Science, 24*(9), 1848–1853.

Fehd, H. M., & Seiffert, A. E. (2008). Eye movements during multiple object tracking: Where do participants look? *Cognition, 108*(1), 201–209.

Henderson, J. M., Brockmole, J. R., & Gajewski, D. A. (2008). Differential detection of global luminance and contrast changes across saccades and flickers during active scene perception. *Vision Research, 48*(1), 16–29. https://doi.org/10.1016/j.visres.2007.10.008

Henderson, J. M., & Hollingworth, A. (1998). Eye movements during scene viewing: An overview. In *Eye Guidance in Reading and Scene Perception* (pp. 269–293). Elsevier.

Henderson, J. M., & Hollingworth, A. (1999). HIGH-LEVEL SCENE PERCEPTION. *Annual Review of Psychology, 50*(1), 243–271. https://doi.org/10.1146/annurev.psych.50.1.243

Henderson, J. M., & Hollingworth, A. (2003). Global transsaccadic change blindness during scene perception. *Psychological Science, 14*(5), 493–497.
Hollingworth, A., & Henderson, J. M. (2000). Semantic informativeness mediates the detection of changes in natural scenes. *Visual Cognition, 7*(1–3), 213–235.
Kolb, H. (2011). *Gross anatomy of the eye*.
Levi, D. M. (2008). Crowding—An essential bottleneck for object recognition: A mini-review. *Vision Research, 48*(5), 635–654.
Magliano, J. P., Miller, J., & Zwaan, R. A. (2001). Indexing space and time in film understanding. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition, 15*(5), 533–545.
Magliano, J. P., & Zacks, J. M. (2011). The impact of continuity editing in narrative film on event segmentation. *Cognitive Science, 35*(8), 1489–1517.
Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*(1), 3–25.
R Core Team. (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/
Rayner, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology, 7*(1), 65–81.
Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin, 124*(3), 372.
Rensink, R. A., O'regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8*(5), 368–373.
Rosenholtz, R. (2016). Capabilities and Limitations of Peripheral Vision. *Annual Review of Vision Science, 2*(1), 437–457. https://doi.org/10.1146/annurev-vision-082114-035733
Scholl, B. J. (2000). Attenuated change blindness for exogenously attended items in a flicker paradigm. *Visual Cognition, 7*(1–3), 377–396.
Simons, D. J. (2000). Current Approaches to Change Blindness. *Visual Cognition, 7*(1–3), 1–15. https://doi.org/10.1080/135062800394658
Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences, 1*(7), 261–267. https://doi.org/10.1016/S1364-6613(97)01080-2
Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences, 9*(1), 16–20.
Smith, T. J., & Henderson, J. M. (2008). Edit Blindness: The relationship between attention and global change blindness in dynamic scenes. *Journal of Eye Movement Research*, *2*(2).
Smith, T. J., Lamont, P., & Henderson, J. M. (2012). The penny drops: Change blindness at fixation. *Perception, 41*(4), 489–492.
Upadhyayula, A., & Flombaum, J. (2020). A model that adopts human fixations explains individual differences in multiple object tracking. *Cognition, 205*, 104418.
Upadhyayula, A., & Henderson, J. M. (2023). Spatiotemporal jump detection during continuous film viewing. *Journal of Vision, 23*(2), 13. https://doi.org/10.1167/jov.23.2.13
Upadhyayula, A., Phillips, I., & Flombaum, J. (2023). Eccentricity advances arrival to visual perception. *Journal of Experimental Psychology: General.* https://doi.org/10.1037/xge0001352
Whitney, D., & Levi, D. M. (2011). Visual crowding: A fundamental limit on conscious perception and object recognition. *Trends in Cognitive Sciences, 15*(4), 160–168.
Wolfe, J. M., & Gray, W. (2007). Guided search 4.0. *Integrated Models of Cognitive Systems*, 99–119.
Zelinsky, G. J. (2001). Eye movements during change detection: Implications for search constraints, memory limitations, and scanning strategies. *Perception & Psychophysics, 63*(2), 209–225.
Zelinsky, G. J., & Neider, M. B. (2008). An eye movement analysis of multiple object tracking in a realistic environment. *Visual Cognition, 16*(5), 553–566.