

Event-Based Warping: A Relative Distortion of Time Within Events

Rui Zhe Goh^{1,2,*}, Hanbei Zhou^{1,*}, Chaz Firestone^{1,2,†}, Ian Phillips^{1,2,†}

¹Department of Psychological & Brain Sciences, Johns Hopkins University

²Department of Philosophy, Johns Hopkins University

*These authors should be considered joint first author

†These authors should be considered joint senior author

Correspondence: rgoh1@jhu.edu, hzhou43@jhu.edu, chaz@jhu.edu, ianbphillips@jhu.edu

Abstract

Objects and events are fundamental units of perception: Objects structure our experience of space, and events structure our experience of time. A striking and counterintuitive finding about object representation is that it can warp perceived space, such that stimuli within an object appear farther apart than stimuli in empty space. Might events influence perceived time in the same way objects influence perceived space? Here, five experiments (N=500 adults) show that they do: Just as stimuli within an *object* are perceived as farther apart in *space*, stimuli within an *event* are perceived as further apart in *time*. Such “Event-based Warping” is elicited both by events characterized by sound (E1), as well as events characterized by silence (E2). Moreover, these effects cannot be explained by surprise, distraction or attentional-cueing (E3 and E4), and also arise cross-modally (from audition to vision; E5). We suggest that object-based warping and event-based warping are both instances of a more general phenomenon in which representations of *structure* — whether in space or in time — generate powerful and analogous relative perceptual distortions.

Public significance statement

Perception segments sensory input received across continuous time into discrete events. For instance, when we listen to a musical piece, our auditory system segments continuous sound waves into the discrete musical notes, phrases, and motifs that we hear. Whereas much previous work has focused on the effects of event segmentation on downstream processes that depend on it (such as attention and memory), here we show that event segmentation can produce an *upstream* effect, distorting our experience of time itself — the very time that was segmented into events in the first place.

Keywords

event segmentation | time perception | multisensory processing

Version

4/11/2025: in press, JEP:General

Transparency and openness

The authors declare no conflicts of interest. This research was supported by NSF BCS #2021053 and a JHU Catalyst Award. This research complies with the Declaration of Helsinki (2023), and received approval from the Homewood Institutional Review Board of Johns Hopkins University.

The hypotheses, methods, and analysis plans for all studies were preregistered prior to data collection. There were no deviations from the preregistration. All study materials, data, and analysis scripts are publicly available, and can be found at <https://osf.io/px5bd>. Readers can also view all experiments as participants experienced them at <https://www.perceptionresearch.org/eventbasedwarping/>.

Author Note

This work was presented at Vision Sciences Society 2023, Society for Philosophy and Psychology 2024, and Association for Scientific Study of Consciousness 2024. We thank the audiences at these conferences for their valuable feedback.

Author Contributions

Rui Zhe Goh, Hanbei Zhou, Chaz Firestone, and Ian Phillips played equal roles in conceptualizing the work, developing the methodology, and editing the manuscript.

Hanbei Zhou played a lead role in investigation, formal analysis, and visualization. Rui Zhe Goh wrote the original manuscript draft, and played a lead role in validation. Chaz Firestone and Ian Phillips supervised the project. Chaz Firestone acquired the funding.

Introduction

Sensory input is a continuous wash of light, sound, and other stimulation. Yet, we do not experience the world as a great blooming, buzzing confusion (James, 1890) but rather as populated by discrete objects and events. Consider the experience of listening to a piece of music: Though the auditory input reaching our sensory organs is a stream of sound waves, we hear individual notes, phrases, transitions, and other structured contents (Phillips et al., 2020). The process responsible for parsing such continuous input into these discrete representations is perceptual event segmentation (Bregman, 1994).

Perceptual event segmentation is a fascinating process in its own right, and previous work has explored how the perceptual system identifies and extracts boundaries corresponding to perceived events (Avrahami & Kareev, 1994; Zacks, 2004; Schapiro et al., 2013). Perceptual event representations also serve as a scaffold for many downstream cognitive processes. Speech comprehension, for instance, relies on discrete event representations of phones and syllables, which must be extracted from continuous speech input (Moore, 2008; Gong et al., 2023). Auditory events also serve as units of auditory attention: Just as we can attend to visual objects in space, so too can we attend to discrete auditory events in time (Shinn-Cunningham, 2008; De Freitas et al., 2014). Moreover, these more basic perceptual event representations (traditionally studied by perception researchers) form the building blocks of more complex event representations (studied by event cognition researchers), which may be further assembled into structures such as scripts, action schemas, and narratives. Such

higher-level event representations are used to model our environmental context (Richmond & Zacks, 2017) and may form the units of storage and retrieval in long-term memory (Huff et al., 2014; Zacks, 2020).

From perception to segmentation and back again

Perceptual event representations clearly lay a foundation for many higher-level cognitive processes. But might event segmentation also exert influence in the other direction—not only supporting the downstream processes that depend on it, but also reaching back upstream to alter perception itself? Since event segmentation relies on discovering and representing the structure of sensory input, it might seem surprising if event representations distorted our perception of the very properties which form the basis of those representations. However, an important clue that this may in fact occur comes from *object* segmentation, a process in which continuous space is carved into discrete, bounded objects. In an illusion known as “Object-based Warping”, two dots within a visual object are perceived as farther apart than equidistant dots in empty space (Vickery & Chun, 2010; Fig 1). This suggests that object segmentation not only guides downstream processes such as attention and working memory, but can also have upstream effects that distort our perception of space—the very space carved up into objects in the first place.

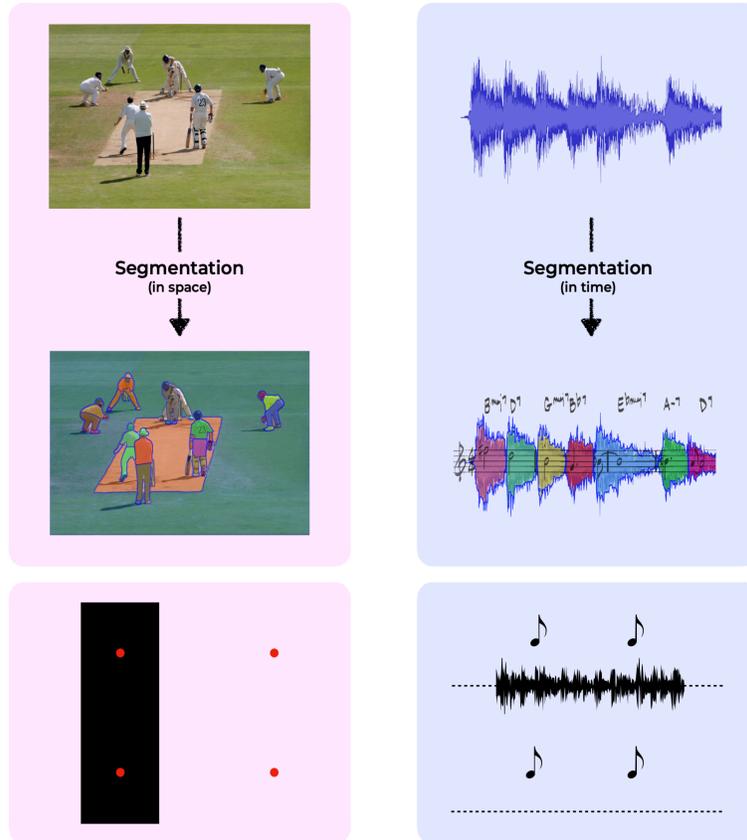


Fig. 1. Perceptual segmentation and segmentation-based illusions. (Top) Perceptual segmentation parses continuous sensory input into discrete representations of objects and events. (Top-left) Continuous visual spatial input is segmented into discrete objects (e.g., people, material objects, regions of the playing field). (Top-right) Continuous temporal input is segmented into discrete events (e.g., musical notes, chords). (Bottom) These segmentation processes can produce perceptual illusions. (Bottom-left) In object-based warping (Vickery & Chun, 2010), two dots in a black rectangle appear farther apart than two dots in empty space. (Bottom right) We propose that spatial object-based warping has a temporal event-based counterpart: Just as two dots within a visual object appear farther apart in *space* than two dots in empty space, two probes within an auditory event are perceived as further apart in *time* than two probes not within an auditory event.

Object-based warping suggests a reciprocal relationship between segmentation and perception: We segment continuous space into discrete, structured elements, which in turn change how that space appears to us. But how general is this relationship? Are previously observed warping phenomena specific to visual objects? Or might they be a feature of perception writ large? Here, we explore this question by asking if an analogous *event*-based warping illusion exists. Just as objects are bounded in space, so events are bounded in *time*. Thus, whereas in object-based warping two dots within an object are perceived to be farther apart in *space*, here we ask whether two probes (e.g., tones or flashes) within an event are perceived as further apart in *time* (Fig. 1). In this way, we examined whether event representations distort our perception of time, just as object representations distort our perception of space.

The present approach

Our experimental strategy was to elicit auditory event representations by presenting subjects with brief soundtracks simulating a novel auditory context (e.g., a busy restaurant). In previous work, this methodology has proven effective in producing event-based illusions (Goh et al., 2023), reflecting the importance of context shifts in triggering event segmentation processes (Radvansky & Copeland, 2006). Here, we used these soundtracks to create discrete auditory events within which two successive probes were presented. Our key question was whether this pair of probes would be perceived as further apart in time when presented within an auditory event (e.g., during a brief period of noise) than when presented outside of an event (e.g., in continuous

silence). We hypothesized that durations within events would be judged as longer than durations outside of events.

Experiment 1: Introducing event-based warping

Do auditory events warp perceived time in ways analogous to object-based warping of perceived space? Experiment 1 tested this question by asking whether tone pairs heard within an event sound further apart in time than tone pairs heard outside of an event (Fig. 2B).

Method

All experiments reported in this paper were pre-registered. An archive of the data, experiment code, stimuli, pre-registrations, and other relevant materials is available at https://osf.io/px5bd/?view_only=86a2de30291c4a6780baf0cedd3185ac. Readers can also view all experiments as participants experienced them at <https://www.perceptionresearch.org/eventbasedwarping/>.

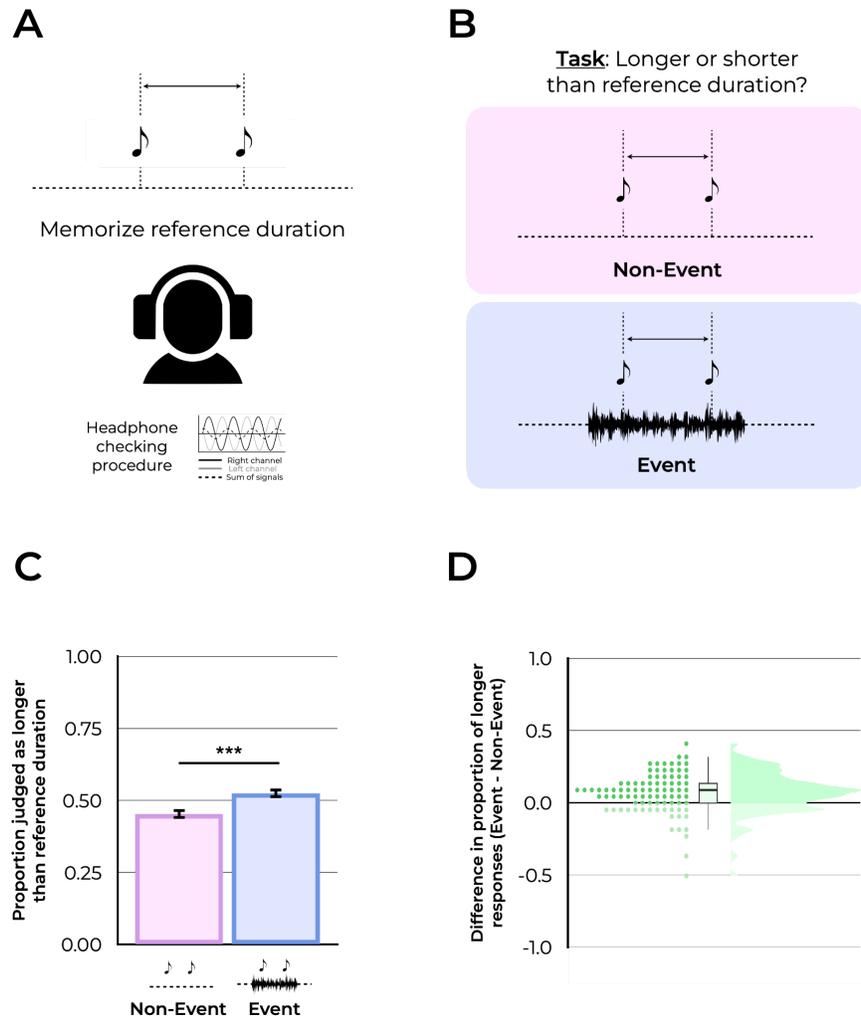


Fig. 2. Stimuli, procedure, and results for Experiment 1. (A) Before each trial, subjects memorized a fixed ‘reference duration,’ later used to make judgments of temporal duration. Subjects were required to wear headphones for the entire duration of the experiment. To ensure compliance, subjects had to complete a headphone screening procedure (which exploits the fact that antiphase tones are heard differently through headphones than loudspeakers; Woods et al., 2017) before being allowed to participate. (B) During each trial, two successive tones were presented and subjects were asked to compare the duration between these tones (known as the ‘test duration’) with the previously memorized reference duration. In Non-Event trials,

tones were presented in complete silence; in Event trials, tones were presented during an auditory event. **(C)** Subjects were more likely to judge a test duration as longer than the reference duration if the test duration occurred within an auditory event. Error bars represent SEM. **(D)** The key effect broken down by subject. Plotted here is the difference in the proportion of “longer” responses for Event trials than for Non-Event trials; a greater difference indicates a stronger effect in the predicted direction.

Participants

We recruited 100 participants from the online platform Prolific. (For validation of the reliability of this subject pool, see Peer et al., 2017.) All subjects provided informed consent and were compensated financially for their participation. No demographic information was collected. The experiments were approved by the Homewood Institutional Review Board of Johns Hopkins University.

Headphone Screening

To promote immersion in the auditory stimuli, subjects were required to wear headphones or earbuds for the entire experiment. To ensure this, all subjects had to pass a headphone screening procedure, which consisted of six trials: On each trial, subjects heard three tones and judged which tone was quietest. One of the three tones was presented 180° out of phase across stereo channels. This tone sounds different on headphones (because each ear receives audio from only one channel) compared to loudspeakers (where both ears receive audio from both channels), making the task easy with headphones but prohibitively difficult with loudspeakers (for details, see

Woods et al., 2017). Subjects had to answer at least five out of six trials correctly before they could participate in the experiment.

Stimuli and Procedure

Our procedure used a version of the classic *reminder task*, in which each trial comprised two presentations, first a standard/reference duration whose magnitude was held fixed throughout the experiment, then a comparison/test duration whose magnitude varied across trials and was compared with the reference duration (for examples of this task, see Lapid et al., 2008; Rammsayer & Ulrich, 2012).

During the instruction phase, subjects heard a pair of reference tones played successively and were asked to memorize the duration between the tones; this is the *reference duration*. Subjects could play the reference tones as many times as it took to memorize the reference duration. The reference duration used throughout a given subject's experimental session was chosen randomly for that subject from three options (1 s, 1.5 s, and 2 s). All tones were 1000 Hz and lasted 200 ms.

Before each trial, subjects had to listen to the reference duration again (they could do this as many times as they wished). There were two types of trials: Event trials and Non-Event trials. Event trials started with 4 s of silence, before an ambient noise onset. The ambient noise was either white noise or the soundscape of a busy restaurant, and was fixed for each subject. After a buffer interval (see below for the possible durations of this interval), two test tones were played successively. The duration between these

test tones is the *test duration*. After an identical buffer interval, the ambient noise offset, followed 2 s later by an on-screen prompt asking subjects to select one of two buttons to indicate if the test duration was longer or shorter than the reference duration.

Non-Event trials were identical to Event trials, except without the ambient noise. These trials started with 4 s of silence followed by a silent buffer interval. Thereafter, two test tones were played successively, followed by another identical silent buffer interval. 2 s later, a prompt appeared on-screen asking subjects to select one of two buttons to indicate if the test duration was longer or shorter than the reference duration.

There were seven test durations in the experiment. The test durations were ratios of the reference duration (0.85, 0.90, 0.95, 1.0, 1.05, 1.10, and 1.15). There were also three possible buffer intervals (0.5 s, 1 s, 1.5 s). Each subject completed 42 experimental trials (7 reference durations x 3 buffer intervals x 2 conditions), presented in random order. In addition, there were two catch trials. In one catch trial, the test duration was twice as long as the reference duration; in the other catch trial, the test duration was half as long as the reference duration. As per our pre-registration, any subject who failed a catch trial was excluded.

For every trial, we also recorded the actual durations of the reference duration and test duration, calculated as the duration between the execution of the command to play the first tone and the execution of the command to play the second tone. This allowed us to calculate a metric for system error, namely the percentage difference between the

actual duration and the intended duration. As specified in our pre-registration, we excluded any subject with at least two trials with a system error larger than 5%.

Results

In total, 11 subjects were excluded in accordance with our pre-registered criteria: 8 due to the criteria mentioned above, and 3 for failing to submit a complete dataset. 89 subjects remained after exclusion.

Subjects judged test durations in Event trials to be longer than test durations in Non-Event trials ($t(88) = 4.47, p < .001$; Fig 2C). This bias was robust to various experimental parameters, arising in each ambient noise condition, and also for each reference duration. These results provide initial evidence that tone pairs within auditory events are judged to be further apart in time than tone pairs not within auditory events — an event-based analog of object-based warping.

Experiment 2: Controlling for the presence of sound

In Experiment 1, tone pairs in Event trials were played in ambient noise, whereas tone pairs in Non-Event trials were played in silence. Could the results of Experiment 1 be explained by the presence of *sound* (perhaps impairing subjects' ability to hear the tones), rather than the presence of event representations? Experiment 2 addressed this confound by 'inverting' the structure of Experiment 1. Instead of hearing events of sound presented in an otherwise silent soundscape, subjects in Experiment 2 were immersed in ambient noise throughout the experiment and experienced brief events of *silence*

interrupting the ambient noise. In Event trials, tone pairs were presented during these silent events, whereas in Non-Event trials, tone pairs were presented while the ambient noise continued playing (Fig. 3A). If the effects in Experiment 1 were caused by interference between the ambient noise and encoding of the tones, then inverting our design in this way should reverse the effects, with test durations in (noisy) Non-Event trials heard as longer than (silent) Event trials. But if the warping effect observed in Experiment 1 was due to the presence of event representations rather than the mere presence of sound, then test durations in (silent) Event trials should still be perceived as longer than test durations in (noisy) Non-Event trials.

Method

The stimuli and procedure for Experiment 2 were identical to Experiment 1 except as specified below.

For the entire duration of the experiment (except for silent periods during Event trials), subjects were immersed in ambient noise. All subjects experienced white noise as the ambient noise.

In Event trials, the ambient noise offset after 4 s, followed by a buffer interval.

Thereafter, two test tones were played successively. After another buffer interval, the ambient noise resumed, followed 2 s later by an on-screen prompt asking subjects to select one of two buttons to indicate if the test duration was longer or shorter than the reference duration.

Non-Event trials were identical to Event trials, except without any offset of the ambient noise, which played continuously. These trials started with 4 s of noise followed by a buffer interval. Thereafter, two test tones were played successively, followed by another buffer interval. 2 s later, a prompt appeared on-screen asking subjects to select one of two buttons to indicate if the test duration was longer or shorter than the reference duration.

To increase statistical power, each subject in Experiment 2 completed 84 experimental trials (7 reference durations x 3 buffer intervals x 2 conditions x 2) and four catch trials, which is twice the number of trials as in Experiment 1. As per our pre-registration, any subject who failed a catch trial was excluded.

Results

11 subjects were excluded in accordance with our pre-registered exclusion criteria: 10 due to the criteria mentioned above, and 1 for failing to submit a complete dataset. 89 subjects remained after exclusion.

Subjects judged test durations in (silent) Event trials to be longer than test durations in (noisy) Non-Event trials ($t(88) = 2.77, p = 0.007$; Fig 3A), demonstrating that event-based warping is not solely caused by effects of sound on temporal experience, and providing further evidence of a general effect of eventhood on perceived duration.

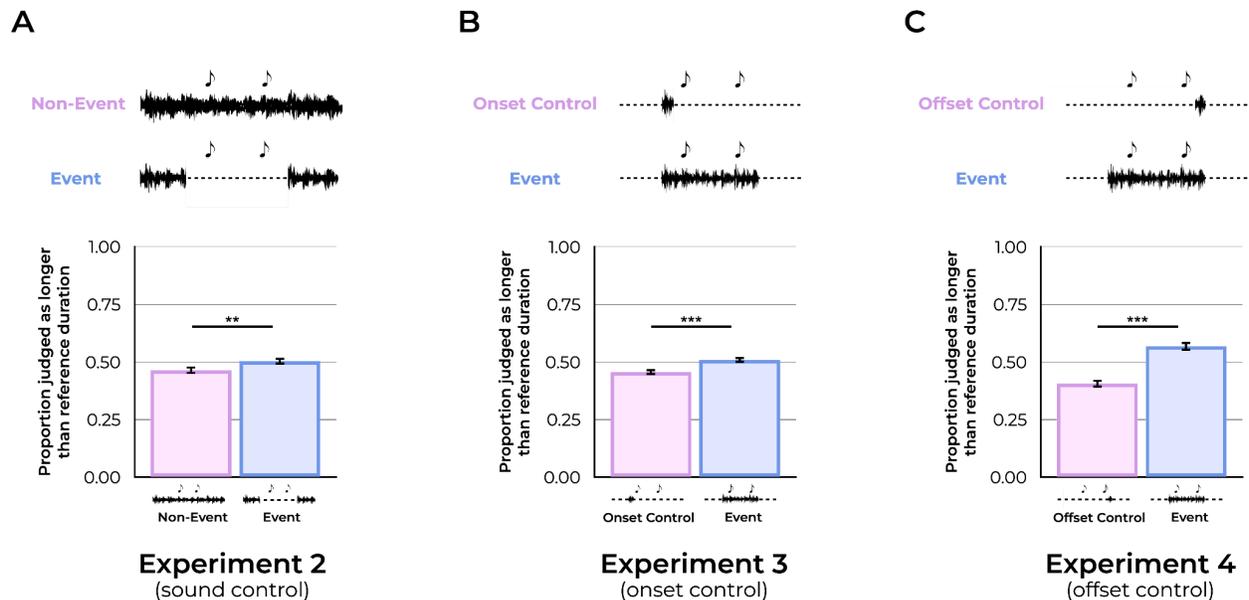


Fig. 3. Stimuli, procedure, and results for Experiments 2–4. (A) To rule out the possibility that event-based warping occurs simply due to the presence of sound, Experiment 2 inverted the structure of Experiment 1. Except as specified, subjects were immersed in ambient noise for the entire experiment. In Non-Event trials, two tones were presented while the ambient noise continued playing. In Event trials, two auditory probes appeared during a brief silence interrupting the ambient noise. Subjects again judged test durations as longer in Event trials as compared to Non-Event trials. (B) To rule out the possibility that event-based warping occurs simply due to sudden changes in sound before tone pair presentation, Experiment 3 modified the Non-Event trials of Experiment 1 to include a brief burst of white noise that onset at the same time as the onset of the auditory event in Event trials. Subjects once again judged test durations as longer in Event trials as compared to Non-Event trials. (C) To rule out the possibility that event-based warping occurs simply due to sudden changes in sound after tone pair presentation, Experiment 4 modified the Non-Event trials of Experiment 1 to include a brief burst of white noise that offset at the same time as the offset of the auditory event in Event trials. Subjects once again judged test durations as longer in Event trials as compared to Non-Event

trials. Error bars represent SEM. See Supplement for results of Experiments 2-4 broken down by subject.

Experiment 3: Controlling for surprise, distraction and cueing

In both Experiments 1 and 2, each tone pair in Event trials was preceded by a sudden change in sound. In Experiment 1, the ambient noise onset just before the tone pair was presented, while in Experiment 2, the ambient noise offset just before tone pair presentation. Could the results of Experiments 1 and 2 be due to subjects being surprised, distracted, or attentionally-cued by these sudden changes in sound?

Experiment 3 controlled for all these different possibilities by modifying the Non-Event trials of Experiment 1 to include a brief burst of white noise before tone pair presentation (Fig. 3B). As a result, both Event trials and Non-Event trials involved sudden changes in sound before tone pair presentation; but in Non-Event trials, the tones were not played *within* an auditory event, since the brief burst of white noise heard against a silent background is likely insufficient to cause a context shift. If event-based warping is solely caused by surprise, distraction, or attentional cueing due to sudden auditory changes before tone pair presentation, then there should be no difference in judgments of test durations between Event trials and Non-Event trials.

Method

The stimuli and procedure for Experiment 3 were identical to Experiment 1 except as specified below.

All subjects experienced white noise as the ambient noise. Non-Event trials were the same as before, except that 4 s after the start of the trial, there was a 250 ms burst of white noise. Consequently, white noise onset at exactly the same time in Event trials and Non-Event trials. As in Experiment 2, each subject completed 84 experimental trials.

Results

13 subjects were excluded in accordance with our pre-registered exclusion criteria: 10 due to the criteria mentioned above, and 3 for failing to submit a complete dataset. 87 subjects remained after exclusion.

Subjects judged test durations in Event trials to be longer than test durations in Non-Event trials ($t(86) = 4.14, p < .001$; Fig 3B). Since both Event trials and Non-Event trials involved a sudden change in sound before the presentation of the tone pairs, this difference in judgments of test durations shows that event-based warping is not solely caused by the effects of surprise, distraction, or attentional cueing due to sudden changes in sound before tone pair presentation.

Experiment 4: Controlling for sound change after tone presentation

In Experiments 1-3, each tone pair in Event trials was followed by a sudden change in sound. In Experiments 1 and 3, the ambient noise offset soon after the tone pair was presented, while in Experiment 2, the ambient noise onset soon after tone pair presentation. Could the results of Experiments 1-3 be due to these sudden changes in

sound interfering with subjects' responses (e.g., by disrupting their working memory representation of the test duration)?

Experiment 4 controlled for this possibility by modifying the Non-Event trials of Experiment 1 to include a brief burst of white noise after tone pair presentation (Fig. 3C). As a result, both Event trials and Non-Event trials involved sudden changes in sound after tone pair presentation; but in Non-Event trials, the tones were not played *within* an auditory event. If event-based warping is solely caused by interference due to sudden auditory changes after tone pair presentation, then there should be no difference in judgments of test durations between Event trials and Non-Event trials.

Method

The stimuli and procedure for Experiment 4 were identical to Experiment 1 except as specified below.

All subjects experienced white noise as the ambient noise. Non-Event trials were the same as before, except that there was a 250 ms burst of white noise after the test tones were presented, which offset at exactly the same time as the offset of ambient noise in Event trials. As in Experiments 2 and 3, each subject completed 84 experimental trials.

Results

17 subjects were excluded in accordance with our pre-registered exclusion criteria: 16 due to the criteria mentioned above, and 2 for failing to submit a complete dataset. 82 subjects remained after exclusion.

Subjects judged test durations in Event trials to be longer than test durations in Non-Event trials ($t(81) = 6.02, p < .001$; Fig 3C). Since both Event trials and Non-Event trials involved a sudden change in sound after the tone pairs were presented, this difference in judgments of test durations shows that event-based warping is not solely caused by interference due to sudden changes in sound after tone pair presentation.

Compared to Experiments 1–3, we observed a non-negligible proportion of subjects with accuracy levels near chance. Thus, we conducted an exploratory follow-up analysis that excluded any subject with an overall accuracy lower than 0.6. The remaining 69 subjects still judged test durations in Event trials to be longer than test durations in Non-Event trials ($t(68) = 6.21, p < .001$; Fig S3B).

Together, the results of Experiments 3 and 4 show that event-based warping occurs even when both Event trials and Non-Event trials involve a sudden change in sound, ruling out the possibility that the effect is solely due to auditory changes interfering with subjects' responses.

Experiment 5: From audition to vision

The previous four experiments demonstrate that tones within an auditory event are heard as further apart in time than tones not within an auditory event. Do auditory events only affect temporal experience within the auditory modality, or does auditory event-based warping affect perception more generally?

Experiment 5 asked whether the warping effects of auditory events can extend beyond the modality they arise in to cause cross-modal warping effects in *vision*. Instead of hearing tone pairs, subjects saw pairs of *visual flashes* presented either within an auditory event or not within an auditory event (Fig. 4A). If auditory events can influence visual time, then flashes within events should be seen as further apart in time than flashes not within an event.

Method

The stimuli and procedure for Experiment 5 were identical to Experiment 1 except as specified below.

All reference durations and test durations were presented as durations between pairs of visual flashes instead of durations between tone pairs. For the entire duration of the experiment, an empty rectangle with a black outline was displayed on the screen (250 px wide; 50 px high; 2 px thick). Each flash was presented by briefly filling this empty rectangle with red (RGB(255, 0, 0)) for 50 ms.

Since the task in this experiment only required subjects to compare durations between visual flashes, subjects could in principle complete the task with the sound off, which would render them insensitive to the auditory events that constitute the main experimental manipulation. To address this potential issue, we included a surprise sound check at the end of the experiment, in which the sound of a word (e.g., “cucumber”) was presented without prior warning, and subjects were asked to indicate which word they heard by choosing from a list of options. As specified in our pre-registration, we excluded any subject who failed the sound check.

Results

8 subjects were excluded in accordance with our pre-registered exclusion criteria mentioned above. 92 subjects remained after exclusion.

Subjects judged test durations in Event trials to be longer than test durations in Non-Event trials ($t(91) = 4.70, p < .001$; Fig 4B), showing that visual flashes perceived as occurring within an auditory event are seen as further apart in time than visual flashes not perceived as occurring within an auditory event. In other words, the warping effects of auditory events extend to the visual modality.

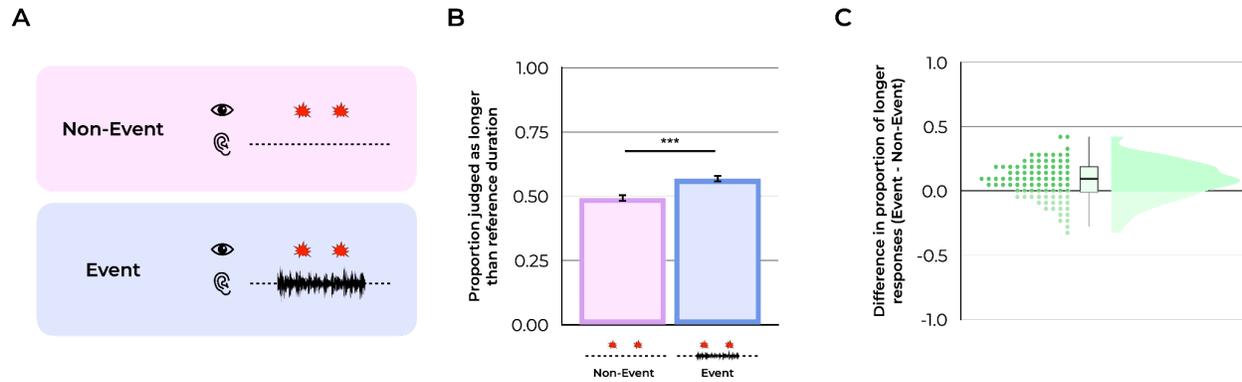


Fig. 4. Stimuli, procedure, and results for Experiment 5. (A) To determine whether the temporal distortion due to auditory events extends to the visual modality, Experiment 5 replaced the tones in Experiment 1 with visual flashes. In Non-Event trials, two flashes were presented in complete silence, whereas in Event trials, two flashes were presented during an auditory event. **(B)** Event-based warping occurred under these cross-modal conditions as well. Error bars represent SEM. **(C)** The key effect broken down by subject. Plotted here is the difference in the proportion of “longer” responses for Event trials than for Non-Event trials; a greater difference indicates a stronger effect in the predicted direction.

General Discussion

Event representations arise as continuous temporal input is parsed into discrete, bounded episodes. Can this segmentation process also exert an ‘upstream’ influence, distorting our perception of time itself? The five experiments reported here introduce a novel temporal illusion which shows that it can: Stimuli occurring within an auditory event are perceived as further apart in time than stimuli occurring outside of an auditory event. This phenomenon occurs not only with events characterized by the presence of

sound (E1), but also with events characterized by silence (E2) — such that the very same acoustic pattern (an interval between test tones played against a silent background in both Non-Event trials in E1 and Event trials in E2) is heard as longer or shorter depending on whether the mind represents that pattern as occurring within an event or not. It also persists when controlling for surprise, distraction and interference due to sudden changes in sound (E3 and E4), and operates across modalities, influencing the perception of visual features (E5). Such *event-based warping* mirrors object-based warping in spatial perception, suggesting that a reciprocal relationship between perception and segmentation constitutes a general feature of perceptual processing.

Relation to other work

The present experiments make a distinctive contribution to a growing body of empirical work investigating how event segmentation influences our perception and memory of temporal properties, such as duration and temporal order. Previous work has focused on comparing responses to stimuli presented *within* an event with responses to stimuli presented *across* events. Such comparison has shown that event boundaries distort judgments of temporal order (DuBrow & Davachi, 2013), that the presence of event boundaries in a sequence can influence memory of its duration (Fenerci et al., 2021; Sherman et al., 2023; Dubrow et al., 2024), and — most relevantly for current purposes — that stimuli across events are perceived as farther apart in time than stimuli within events. As an example of this third finding, Ongchoco and Scholl (2023) had subjects reproduce a previously experienced rhythm, and found that the duration between two

beats was reproduced as longer when the two beats were heard across an event boundary compared to when the two beats were heard within the same event, suggesting that temporal experience is dilated across event boundaries (see also Bangert et al., 2020).

In contrast to previous work, we compared perception of stimuli *within* an event to perception of stimuli *not within an event* — thus uncovering more basic effects of event representation on perception. Our findings can be integrated with Ongchoco and Scholl's findings by positing a hierarchy of temporal dilations, such that temporal experience is dilated within events relative to temporal experience not within an event, but temporal experience is further dilated across events relative to temporal experience within events.

The reciprocal relationship between auditory event representations and temporal experience found here mirrors the reciprocal relationship between visual object representations and spatial experience demonstrated in object-based warping (Vickery & Chun, 2010), suggesting that these warping phenomena are not restricted to a specific sense modality or stimulus set, but rather constitute a fundamental feature of perception more generally. The analogy between object-based warping and event-based warping may seem surprising given the differences in the mechanistic requirements of object-based and event-based tasks. For instance, our experiments required the use of working memory to compare the reference duration with the test duration, whereas working memory is not obviously required in tasks used to

demonstrate object-based warping, in which all stimuli are presented simultaneously (for more on the role of working memory in event representation, see Lu et al., 2019 and Gu et al., 2020).

Nonetheless, our findings pattern with a growing literature reporting event-based counterparts to object-based phenomena. For instance, De Freitas et al. (2014) discovered an event-based analog of object-based attention (Egley et al., 1994). In their paradigm, an auditory cue facilitated detection of a subsequent probe more effectively when the cue and probe were presented in the same event as compared to when cue and probe were presented in different events, suggesting that attention ‘spreads’ within an event, just as it spreads within an object (for other work making empirical connections between objects and events, see Ongchoco & Scholl, 2022; Yousif & Scholl, 2019; Lee et al., 2024; Papafragou & Ji, 2023; Wellwood et al., 2018). Whereas previous work on object-event analogies has focused on downstream effects, here we show that the analogy extends to upstream influences of object and event representations on spatial and temporal experience respectively. The fact that analogous phenomena occur with both events and objects suggests that event segmentation and object segmentation recruit mechanisms with similar underlying principles, even if they might be implemented by different computational systems or in different brain areas. We suggest that the mechanisms underlying object and event segmentation should be understood in light of a common goal that both processes share, namely to recover information about discrete entities in the external world from continuous sensory input.

Constraints on Generality and Future Directions

Our experiments demonstrate that event-based warping generalizes across a variety of different experimental conditions, such as with different auditory stimuli (e.g., restaurant and white noise), with different reference durations, and with both visual and auditory probes. Future work could explore other boundary conditions. For instance, to establish the basic effect, we focused on using simple auditory events demarcated by onsets and offsets of sound. Does event-based warping occur with more complex auditory events, such as musical phrases, or events in other sense modalities? Our focus was also on significantly extended perceptual events of between one and two seconds within which other things could be perceived as occurring. Future work could explore whether the same pattern of results holds at significantly shorter and longer timescales, as well as the effects on time perception of ‘punctate’ events such as our auditory probes (see Yates et al., 2023 for more on the distinction between punctate and extended events — ‘moments’ and ‘periods’ in their terminology). Finally, our finding that the relative temporal distortion elicited by auditory events generalizes to vision (E5) raises the question of whether time experienced in other modalities (e.g., touch, smell) would also be affected by auditory event-based warping.

Future work should also further explore the mechanisms underlying warping effects of the kind we observe here. One possibility, congruent with findings that attention spreads within event representations, is that stimuli represented as bound to an event representation (such as two tones occurring within an event) are subject to deeper

processing, which in turn produces the experience of more time passing between the stimuli. Object- and event-based warping may thus be evidence of deeper processing dedicated to the contents of object and event representations alike. The causes and boundary conditions of object-based warping are still unknown (e.g., Lebed et al., 2023; Baker et al., 2024). It thus seems promising to investigate both event- and object-based warping in tandem, since as discussed the two phenomena likely stem from related sources.

A further, related question concerns how far the analogy between events and objects extends. The perception of objects has been subject to decades of extensive empirical study (Spelke, 1990; Ullman, 2000; Scholl, 2001). Thus, a fruitful research program could continue to explore whether other object-based phenomena have event-based analogs (De Freitas et al., 2014). For instance, does object-based spatial neglect (Tipper & Behrmann, 1996; Walker, 1995) have an event-based counterpart? Does the object-specific preview benefit (Kahneman et al., 1992) have an analog in event representation? These explorations could reveal the extent to which object and event segmentation share common mechanistic principles. More generally, we hope that our findings inspire future work investigating the similarities and differences between event segmentation and object segmentation.

References

Avrahami, J., & Kareev, Y. (1994). The emergence of events. *Cognition*, 53(3), 239-261.

Baker, N., Vushaj, A., Friebe, W., & Yahng, J. (2024). Is object-based warping solely object-based? *Journal of Experimental Psychology: Human Perception and Performance*, 50(6), 605–625.

Bangert, A. S., Kurby, C. A., Hughes, A. S., & Carrasco, O. (2020). Crossing event boundaries changes prospective perceptions of temporal length and proximity. *Attention, Perception, & Psychophysics*, 82, 1459–1472.

Bregman, A. S. (1994). *Auditory scene analysis: The perceptual organization of sound*. MIT press.

De Freitas, J., Liverence, B. M., & Scholl, B. J. (2014). Attentional rhythm: A temporal analogue of object-based attention. *Journal of Experimental Psychology: General*, 143(1), 71–76.

DuBrow, S., & Davachi, L. (2013). The influence of context boundaries on memory for the sequential order of events. *Journal of Experimental Psychology: General*, 142(4), 1277–1286.

Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, 123(2), 161–177.

Fenerci, C., da Silva Castanheira, K., LoParco, M., & Sheldon, S. (2021). Changes in the experience of time: The impact of spatial information on the perception and memory of duration. *Quarterly Journal of Experimental Psychology*, 74(3), 471–482.

Goh, R. Z., Phillips, I. B., & Firestone, C. (2023). The perception of silence. *Proceedings of the National Academy of Sciences*, 120(29), e2301463120.

Gong, X. L., Huth, A. G., Deniz, F., Johnson, K., Gallant, J. L., & Theunissen, F. E. (2023). Phonemic segmentation of narrative speech in human cerebral cortex. *Nature Communications*, 14(1), 4309.

Gu, Q., Wan, X., Ma, H., Lu, X., Guo, Y., Shen, M., & Gao, Z. (2020). Event-based encoding of biological motion and location in visual working memory. *Quarterly Journal of Experimental Psychology*, 73(8), 1261–1277.

Huff, M., Meitz, T. G., & Papenmeier, F. (2014). Changes in situation models modulate processes of event perception in audiovisual narratives. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(5), 1377–1388.

James, W. (1890). *The Principles of Psychology*. Henry Holt and Company.

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24(2), 175–219.

Lapid, E., Ulrich, R., & Rammsayer, T. (2008). On estimating the difference limen in duration discrimination tasks: A comparison of the 2AFC and the reminder task. *Perception & psychophysics*, 70(2), 291–305.

Lebed, A., Scanlon, C., & Vickery, T. J. (2023). Expansion and compression of space within and beyond the boundaries of an object. *Attention, Perception, & Psychophysics*, 85(2), 387–403.

Lee, S. H. Y., Ji, Y., & Papafragou, A. (2024). Signatures of individuation across objects and events. *Journal of Experimental Psychology: General*, 153(8), 1997–2012.

Lu, X., Ma, X., Zhao, Y., Gao, Z., & Shen, M. (2019). Retaining event files in working memory requires extra object-based attention than the constituent elements. *Quarterly Journal of Experimental Psychology*, 72(9), 2225–2239.

Moore, B. C. (2008). Basic auditory processes involved in the analysis of speech sounds. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1493), 947–963.

Ongchoco, J. D. K., & Scholl, B. J. (2022). Scaffolded attention in time: 'Everyday hallucinations' of rhythmic patterns from regular auditory beats. *Attention, Perception, & Psychophysics*, 84, 332–340.

Ongchoco, J. D. K., Yates, T. S., & Scholl, B. J. (2023). Event segmentation structures temporal experience: Simultaneous dilation and contraction in rhythmic reproductions. *Journal of Experimental Psychology: General*, 152(11), 3266–3276.

Papafragou, A., & Ji, Y. (2023). Events and objects are similar cognitive entities. *Cognitive Psychology*, 143, 101573.

Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153–163.

Phillips, M., Stewart, A. J., Wilcoxson, J. M., Jones, L. A., Howard, E., Willcox, P., ... & De Roure, D. (2020). What determines the perception of segmentation in contemporary music? *Frontiers in Psychology*, 11, 1001.

Radvansky, G. A., & Copeland, D. E. (2006). Walking through doorways causes forgetting: Situation models and experienced space. *Memory & Cognition*, 34, 1150–1156.

Rammsayer, T., & Ulrich, R. (2012). The greater temporal acuity in the reminder task than in the 2AFC task is independent of standard duration and sensory modality.

Canadian Journal of Experimental Psychology, 66(1), 26-31.

Richmond, L. L., & Zacks, J. M. (2017). Constructing experience: Event models from perception to action. *Trends in Cognitive Sciences*, 21(12), 962-980.

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure.

Nature Neuroscience, 16(4), 486-492.

Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80(1-2), 1-46.

Shinn-Cunningham, B. G. (2008). Object-based auditory and visual attention. *Trends in Cognitive Sciences*, 12(5), 182-186.

Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, 14(1), 29-56.

Tipper, S. P., & Behrmann, M. (1996). Object-centered not scene-based visual neglect.

Journal of Experimental Psychology: Human Perception and Performance, 22(5), 1261-1278.

Ullman, S. (2000). *High-level vision: Object recognition and visual cognition*. MIT press.

Vickery, T. J., & Chun, M. M. (2010). Object-based warping: An illusory distortion of space within objects. *Psychological Science*, 21(12), 1759-1764.

Walker, R. (1995). Spatial and object-based neglect. *Neurocase*, 1(4), 371-383.

Wellwood, A., Hespos, S. J., & Rips, L. (2018). How similar are objects and events?. *Acta Linguistica Academica*, 65(2-3), 473-501.

Woods, K. J., Siegel, M. H., Traer, J., & McDermott, J. H. (2017). Headphone screening to facilitate web-based auditory experiments. *Attention, Perception, & Psychophysics*, 79, 2064-2072.

Yates, T. S., Sherman, B. E., & Yousif, S. R. (2023). More than a moment: What does it mean to call something an 'event'?. *Psychonomic Bulletin & Review*, 30(6), 2067-2082.

Yousif, S. R., & Scholl, B. J. (2019). The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time. *Cognition*, 185, 121-130.

Zacks, J. M. (2004). Using movement and intentions to understand simple events. *Cognitive Science*, 28(6), 979-1008.

Zacks, J. M. (2020). Event perception and memory. *Annual Review of Psychology*, 71, 165-191.