

Tuning to the Task at Hand: Processing Goals Shape Adults' Attention to Unfolding Activity

Jessica E. Kosie (jkosie@uoregon.edu)

Department of Psychology
University of Oregon, Eugene, OR 97403, USA

Dare A. Baldwin (baldwin@uoregon.edu)

Department of Psychology
University of Oregon, Eugene, OR 97403, USA

Abstract

Human activity generates dynamic, multi-modal sensory streams. Effectively processing this complex flow of information on-the-fly is essential if one is to remember and respond to others' action, anticipate what they might do next, and learn how to perform new actions. Selectively attending to information-rich regions of activity seems key to fluent processing. However, what counts as information-rich likely depends on numerous factors including relevance to the causal structure of the activity, local opportunity for repeated viewing, and processing goals of the observer. We explored the influence of these factors on observers' attention to a dynamic, novel activity sequence. A performance context elicited nuanced differences in processing in contrast to a *remember* context. Specifically, individuals given a *perform* context tuned in to causally distinct regions of the action stream and fine-level event details. These findings provide altogether new information regarding how processing rapidly reorganizes around novel activity and responds to the processing task at hand.

Keywords: event processing; action segmentation; context effects

Consider the activity of knitting. No doubt you've heard of this activity, can recognize someone who is engaging in it, and at least globally understand the actor's goals (i.e., transforming a strand of yarn via a complex, repetitive, and very extended sequence of stitches into a piece of patterned material). For those who aren't knitters, the details are opaque, but the general gist of the activity is understood. But what if a non-knitter opted to acquire this skill? Suddenly knitting behavior would be processed in a very different way, presumably with a focus on discovering what motions are actually needed to transform yarn into garments. Precisely what are these changes in processing that the learner initiates? The present research took steps toward answering this question. Finding such answers is foundational to cognitive science, shedding light on basic processes that make observational learning possible, with the hope of ultimately enabling us to assist those for whom observational learning is difficult or disrupted. We begin by reviewing current research on event processing and then present a novel study in which we explore the influence of observers' processing goals on their online attention to unfolding activity.

Segmentation is Key to Fluent Event Processing

Most research in the domain of event processing focuses on activities that are at least moderately familiar to observers. This body of research suggests that, in processing familiar activity, observers chunk unfolding sensory streams into discrete units that are demarcated by event boundaries. Typically, these event boundaries coincide with transitions between one unit of action and another; for example, the moment at which an actor's hand contacts the handle of a mug when reaching for a cup of coffee. Observers overwhelmingly agree when asked to explicitly identify the location of such boundaries in unfolding activity sequences (e.g., Newton, 1973; Zacks, Tversky, & Iyer, 2001; Kurby & Zacks, 2008) as well as when they are asked to scale their segmentation judgments up or down in terms of the grain at which they identify boundaries (e.g., Zacks & Swallow, 2007). To illustrate this granularity, the above-mentioned mug-grasping event might represent a fine-level event boundary in a coffee-making event sequence. In the same action sequence, the "coffee making" event might begin with a coarse boundary at which an actor, having just entered the kitchen, removes a bag of grounds from the cupboard. The coarse boundary demarcating the end of the event sequence – and representing structure within the sequence at more of a "gist" level – might occur once the actor has finished pouring herself a cup of coffee and replaces the coffee pot.

Across a variety of implicit probes of processing, including behavioral tasks, fMRI, and pupillometry, researchers have demonstrated that such targeting of event boundaries occurs automatically as observers view unfolding event sequences (Newton & Engquist, 1976; Schwan & Garsoffky, 2004; Zacks et al., 2001; Tanaka & Baldwin, in preparation). Recently, Hard, Recchia, and Tversky (2011) demonstrated that observers advancing at their own pace through a slideshow composed of frames extracted at a regular increment (e.g., 500 msec) from streaming activity "dwell" longer on slides depicting event boundaries relative to slides depicting within-event content. Further, the amount of time spent dwelling on slides directly corresponds to the level of event hierarchy represented by the slide content. That is, viewers dwell longest on slides representing coarse-level boundaries, dwell less on slides that occur at fine-level

boundaries, and dwell least on slides representing non-boundary content.

It is also worth noting that viewers' sensitivity to event boundaries – whether measured via explicit judgments or implicit measures such as dwell time – predicts other aspects of their event processing, such as their memory for event sequences and the ability to enact such sequences themselves (Kurby & Zacks, 2011; Sargent et al., 2013; Zacks, Speer, Vettel, & Jacoby, 2006; Bailey, Kurby, Giovanetti, & Zacks, 2013; Hard et al., 2011). Thus skill at detecting event units as activity unfolds across time is relevant to memory for that activity and the ability to enact it oneself, which are both hallmarks of observational learning.

Event Boundaries are Low-Predictability Regions

Event boundaries thus appear to be moments in the event stream that are key to observers' fluent processing of unfolding activity sequences. A current conception of the role that event boundaries play in processing is that they represent information-rich regions of the event stream precisely because event boundaries are points that coincide with reductions in the ability to predict what will happen next as activity unfolds (Kosie & Baldwin, 2016; Kurby & Zacks, 2008; Ross & Baldwin, 2015; Zacks, Kurby, Eisenberg, & Haroutunian, 2011). For example, once a reach for a coffee mug has been detected, much is highly predictable, at least until the coffee mug has been grasped. At this juncture, the actor could pursue any number of subsequent acts. She might bring the mug to her mouth to take a drink, move the mug to the sink to wash it, hand the mug to a friend, and so forth. Attending to activity at this low predictability juncture would enable observers to glean important new information about what occurs next. Increasing attention at event boundaries therefore enables observers to gather vital information that guides subsequent processing.

When activity is familiar, this account seems highly intuitive. However, to return to our initial example, if one has little understanding of the actual mechanics of knitting, the ability to target event boundaries seems a significantly more difficult task. On first viewing of highly novel activity, it is unlikely that one can efficiently target event boundaries as almost everything is low predictability and thus information-rich. Recently, Kosie and Baldwin (2016; under revision) used the dwell-time paradigm to demonstrate that observers' processing of novel events reconfigures rapidly. Observers were instructed to advance at their own pace through slideshows depicting novel and familiar methods of shoelace tying. On first viewing of the novel method, observers failed to show the typical systematic increase in attention to event boundaries. But by the second viewing of the novel tying method, observers elevated attention to event boundaries, indicating rapid boundary identification and consequent reorganization of attention to favor boundaries. In addition, novice observers tended to linger on the causally distinctive regions of the event stream that were particularly important for performing the novel activity, in this case the features of the shoelace tying event differentiating novel from familiar

methods. Perhaps increasing attention to causally distinctive regions of the event stream enabled observers to extract the fine-level structure important for carrying out the novel activity.

Importantly, in the research just described, observers were given no instruction to guide their event viewing other than to advance at their own pace through the unfolding activity. In particular, they were given no guidance about how to attend to the shoelace tying event, or to what purpose. Upon completion of this study, however, a subset of participants were told that they would have the opportunity to learn to enact the novel method of shoelace tying. As part of this pilot study, these participants were invited to advance once more through the slideshow depicting the novel activity and then were given the opportunity to try the method themselves. When participants advanced through the slideshow with the goal of learning to perform the actions themselves, their dwell times increased substantially relative to the dwelling they had displayed on their previous, uninstructed, viewing. Further, these increases in dwell time were especially pronounced in relation to causally distinctive regions of activity. However, these pilot findings don't clarify expressly why such dwell-time changes emerged; for example, it is unclear whether an enactment goal specifically generated change in dwell-time patterns, or whether any guidance in how to pay attention would elicit similar alteration.

Prior research supports the notion that context markedly influences processing of event sequences. For example, in their change blindness research Simons and Chabris (1999) famously showed that, when given instructions that focused attention on detail within an activity stream, many participants utterly failed to notice a man in a gorilla suit traipse past in a video of unfolding activity. In contrast, the man in the gorilla suit was readily noticed by the vast majority when the context emphasized more global processing of the activity. Especially relevant to the issues of specific interest here, Blakemore and Decety (2001) reported that cortical activation patterns detected in fMRI differed when participants watched an activity sequence with the instruction to later perform it, relative to the instruction to remember it. However, details about what changes in terms of processing during perform versus remember contexts remain unclear. The dwell-time paradigm offers a potential window on the details of such processing differences.

Overview of the Current Study

We employed the dwell-time paradigm to investigate the extent to which instructions to remember versus perform yielded differences in observers' processing of a novel activity sequence. Participants were asked to view the slideshow used in Kosie and Baldwin (2016; under revision) that depicted a novel shoelace tying sequence. Before participants began, half were instructed to watch the activity so that they could later perform it themselves while the other half of participants were instructed to watch so that they could later remember it. Participants then used a computer mouse to advance at their own pace through the slideshow, in

which the actor first tied her right shoe and then her left shoe (thus they had two viewings of the novel activity). After all participants had advanced through the slideshow they were asked to demonstrate, on a wooden shoe with laces, the method of shoelace tying that they had just viewed.

We anticipated a lower mean per-slide dwell time for *remember* than *perform* instructions. Relative to *perform* instructions, the *remember* processing context was expected to yield dwell-time patterns more like those observed in earlier research in which instructions were to simply watch the novel activity. In replication of our prior research, we expected that, when given *remember* instructions: a) observers would attend longer to event boundaries, with attention particularly enhanced to coarse-grain level boundaries, and b) dwell times would be particularly elevated to distinctive regions of the activity, and especially on second viewing. However, we also predicted that *perform* instructions might elicit increased attention to causally distinctive content, perhaps resulting in higher dwell times to fine-level event boundaries and distinctive regions.

Method

Stimuli

In the current study, participants viewed only one slideshow (depicting the novel method of shoelace tying used in Kosie & Baldwin, 2016; under revision). This novel *twist* method of shoelace tying involved the actor making an initial knot, slipping her pinky fingers under the laces, making a pincer grasp, twisting the pincers around to meet in the middle, grabbing the laces, pulling them through to create a bow, and double-knotting the shoe. A video depicting this novel method of shoelace tying can be viewed at: <https://osf.io/8rpkf/>. The slideshow was created by extracting one still frame every second from a 115-second video of an actor demonstrating the *twist* method (the resulting slideshow thus consisted of 115 unique slides). The extraction rate of one frame per second is consistent with prior research using the dwell-time paradigm (e.g., Hard et al., 2011).

Slide Classification

Two expert coders, with extensive experience in event processing research, first defined regions of the slideshow as causally distinctive (e.g., unique to the *twist* method of shoelace tying) versus non-distinctive (e.g., common to any method of shoelace tying, such as the initial knot at the beginning and double-knot at the end). Individual slides were then classified as depicting boundaries or within-unit content, and boundary slides were further classified at the coarse-grained or fine-grained level of hierarchical structure. These judgments were validated by a sample of naïve research participants. As is typical of naturalistic activity, the precise number of slides falling into the distinctive / non-distinctive

and coarse / fine / within categories differed across viewings (see Table 1). The slide classification process is described in further detail in Kosie & Baldwin, 2016 and Kosie & Baldwin, under revision.

Table 1: Number of slides at each level of structure across viewings and for distinctive and non-distinctive regions.

	Distinctive		Non-Distinctive		Total
	First Viewing	Second Viewing	First Viewing	Second Viewing	
Coarse	0	0	5	6	11
Fine	5	5	8	9	27
Within	15	16	25	20	76
Total	20	21	38	35	114

Note: One slide was classified as the “switch” from first to second viewing and is thus not included in these values.

Participants and Procedure

130 undergraduates (69% female, $M_{age} = 19$ years) participated in exchange for course credit. During an initial phase designed to familiarize participants with the self-paced slideshow format, participants advanced at their own pace through two brief slideshows unrelated to shoelace tying. They were then told that they would use the computer mouse to advance at their own pace through another slideshow (the shoelace tying activity). At this juncture, participants were given either *remember* or *perform* instructions. Participants in the *remember* condition were told: “*You will later be tested for your ability to remember the action that occurred in the slideshow. Please watch the slideshow so that you can remember the action later.*” Participants in the *perform* condition were given the exact same instructions, but with the word *perform* instead of *remember*. The inclusion of either the word *remember* or *perform* in these instructions was the sole difference between the two conditions; participants’ experiences were otherwise identical.

After hearing these instructions, participants advanced at their own pace through the novel method of shoelace tying¹. Participants’ dwell times, or latency between mouse clicks from one slide to the next, were recorded using PsychoPy (Pierce, 2007), a user-friendly experimental control system written in Python. After slideshow viewing, regardless of condition, participants were handed a wooden shoe and asked to demonstrate the novel method of shoelace tying.

Results

Data Preparation

Raw dwell times were subjected to the standard treatment for dwell-time data (i.e., Hard et al., 2011; Kosie & Baldwin, 2016; under review). First, as is typical of reaction-time data,

¹ Data from a subsequent repetition of this entire task are not included in the current analyses but will be reported in a later manuscript.

dwell times (in milliseconds) were log transformed to remove positive skew. Next, outlying dwell times (> 3 SD above the group mean) were removed. Individual participants were excluded if more than 10% of their dwell times met this criterion, resulting in the exclusion of one participant. To account for participants' tendency to speed up as slideshows progress, data from the remaining 129 participants (63 receiving *remember* and 66 receiving *perform* instructions) were individually fitted to a power function. Residuals from these power functions were used as the dependent-variable in analyses targeting within-subjects effects (i.e., boundary vs. within; distinctive vs. non-distinctive).

When necessary, a Greenhouse-Geisser correction was applied to degrees of freedom to address sphericity violations. To adjust for multiple comparisons, Bonferroni correction was applied to all *post hoc* pairwise comparisons.

No Global Dwelling Increase for *Perform* Versus *Remember*

We first examined the overall influence of instructions to remember versus perform on observers' processing of the novel event stream. In this analysis, we simply asked whether mean per-slide dwell time differed for slideshows in which participants were instructed that they would later be asked to remember or later asked to perform the novel shoelace tying method. For these analyses, we used participants' log10 dwell times as the process of residualization substantially attenuates overall group differences. Mean per-slide log10 dwell time for participants receiving *remember* instructions ($M = 2.71, SD = 0.19$) did not significantly differ from that of participants who received *perform* instructions ($M = 2.73, SD = 0.20$), $t(127) = -0.68, p = 0.50, d = -0.12, 95\%CI[-0.09, 0.04]$. Though the instructions to *remember* or *perform* did not differentially affect participants' average per-slide dwell time, mean per-slide dwell times across both levels of instruction ($M = 2.72, SD = 0.19$) were significantly higher than dwell times to the same novel activity in previous research (in which instructions were simply to observe the activity sequence) ($M = 2.53, SD = 0.19$), $t(260) = -8.15, p < .001, d = -1.01, 95\%CI[-0.24, -0.15]$. Across the two studies, it seems that processing instructions in general elicit increased overall dwelling, but instructions to *remember* or *perform* do not differentially influence this global increase in dwelling. We next turned to exploring the effects of instructions on more nuanced facets of processing.

Individuals Given *Perform* Instructions Target Distinctive Content

In our next analysis, we examined the extent to which attention to the distinctive versus non-distinctive regions of the activity sequence differed with respect to instructions and viewing. The analysis of interest here was a 2 (Instructions: Remember vs. Perform) x 2 (Viewing: First vs. Second) x 2 (Region: Distinctive vs. Non-Distinctive) ANOVA with instructions varying between subjects and viewing and region varying within-subjects. Our dependent-variable for this analysis was average residualized dwell time.

Most notably, we found a significant interaction between instructions and slideshow region, $F(1, 127) = 8.28, p = .004, \eta_p^2 = .06$, illustrated in Figure 1 (with means and standard deviations). There was no main effect of viewing, nor did viewing interact with instructions or region, $ps > .06, \eta_p^2 < .02$. Simple-effects comparisons exploring the locus of the interaction between instructions and region, revealed that perform instructions elicited a significant elevation in dwell times to the distinctive relative to non-distinctive region $p = .05, d = -0.60$, whereas remember instructions did not do so, $p = .63, d = 0.35$. Furthermore, dwelling to the distinctive region was significantly elevated by *perform* relative to *remember* instructions, $p = .02, d = 0.51$, whereas dwelling to the non-distinctive region was significantly reduced by *perform* versus *remember* instructions, $p = .02, d = 0.50$.

In sum, individuals given *remember* instructions did not differentiate distinctive versus non-distinctive event content in relation to how they deployed their attention. However, *perform* instructions elicited selectively elevated attention to distinctive event content while reducing attention to non-distinctive content. Together, these findings indicate that participants responded to the perform instruction by tuning in to distinctive content, while ignoring information not relevant to the task at hand (i.e., to learn to enact the novel method of shoelace tying). Also noteworthy was that these patterns emerged to an equivalent degree across first and second viewings. None of the effects involving viewing were significant, suggesting that, when given *perform* instructions, an increase in attention to distinctive regions occurs on first viewing of novel activity and remains stable.




Figure 1. Average residualized dwell times (+/- SE) to distinctive and non-distinctive regions across *remember* and *perform* instructions. Means and SDs reported below bars.

Context Influenced Granularity of Processing

Also of particular interest was the extent to which instructions influenced attention to slides at varying levels of hierarchical structure and how these effects changed across repeated viewing. Thus, the goals of our next set of analyses were threefold: 1) examine replication of previous dwell-time patterns (i.e., boundary and hierarchical advantage effects),

2) explore the extent to which these effects differed across regions of the novel activity sequence, and 3) investigate the influence of instructions to *remember* versus *perform* on these attentional patterns. Because coarse-level boundaries occur only in non-distinctive portions of the event stream, the region factor could not be included in this analysis.

A 2 (Instructions: Remember vs. Perform) x 2 (Viewing: First vs. Second) x 3 (Slide Type: Coarse vs. Fine vs. Within) mixed-design ANOVA (with average residualized dwell times as the dependent variable) revealed no significant main effects of instructions or viewing, $ps > .30$. However, a significant main effect of slide type emerged, $F(1.49, 188.87) = 15.07, p < .001, \eta_p^2 = .11$, replicating previous dwell-time research. A more focused examination of this main effect revealed that average residualized dwell times were longer to boundary than within-unit slides at both the coarse, $p < .001, d = 0.37$, and fine, $p < .001, d = 0.62$, levels of structure, but coarse and fine level boundaries did not significantly differ from one another, $p = .49, d = 0.13$. Thus, averaging across instruction type and viewing, observers exhibited a boundary advantage (replicating previous research using the dwell time paradigm), but average dwell times did not differ with respect to hierarchical structure (failing to replicate the previously observed hierarchical advantage).

However, dwell times related to the slide type variable interacted with both instructions, $F(1.49, 188.87) = 8.23, p = .001, \eta_p^2 = .06$, and viewing, $F(1.60, 203.94) = 7.16, p = .002, \eta_p^2 = .05$. These effects are depicted in Figure 2, which includes means and standard deviations. For those who received the *remember* instructions, the effect of slide type changed significantly across viewing, $F(1.56, 96.74) = 6.78, p = .004, \eta_p^2 = .10$. Specifically, on first viewing, dwell times to coarse, fine, and within-unit slides did not differ, $ps > .71, ds < 0.35$. By the second viewing, however, dwell times exhibited the predicted hierarchical linear trend; greater to coarse-level boundaries than both fine-level boundaries, $p < .001, d = 0.80$, and within-unit slides, $p < .001, d = 0.96$, and dwell times to fine-level boundaries greater than dwell times to within-unit slides, $p = .003, d = 0.85$. Also, for participants receiving the *remember* instructions, dwell times to coarse level boundaries increased across first and second viewing (though this increase was not statistically significant when controlling for multiple comparisons, $p = .16, d = 0.43$). Dwell times to fine-level boundaries did not differ across viewing, $p = .96, d = 0.29$, while dwell times to within-unit slides decreased, $p = .01, d = 0.59$. Conversely, for participants who received the *perform* instructions, the effect of slide type was significant, $F(1.61, 104.64) = 7.90, p = .001, \eta_p^2 = .11$, but did not interact with viewing $F(2, 122) = 2.12, p = .12, \eta_p^2 = .03$. Further analyses of the main effect of slide type revealed that dwell times to coarse-level boundaries were shorter than dwell times to fine-level boundaries, $p = .47, d = -0.18$, and longer than dwell times to within-unit slides, $p = 0.05, d = 0.25$, though neither of these effects were significant. The locus of this effect thus seemed to be that, across both viewings, dwell times to fine-level boundaries

were significantly longer than dwell times to within-unit slides, $p < .001, d = 0.87$.

Taken together, these results indicate a marked contrast in participants' attentional profile in relation to *remember* versus *perform* instructions: when participants were told they would later have to remember the novel activity, they progressively increased attention to the coarse level of structure across viewings. In contrast, participants who were told they would later have to perform the novel activity attended to the fine level of structure across viewings. That is, it appears that instructions to perform elicit processing of the details required to successfully perform this novel method of shoelace tying, while instructions to remember elicited more gist-level processing.

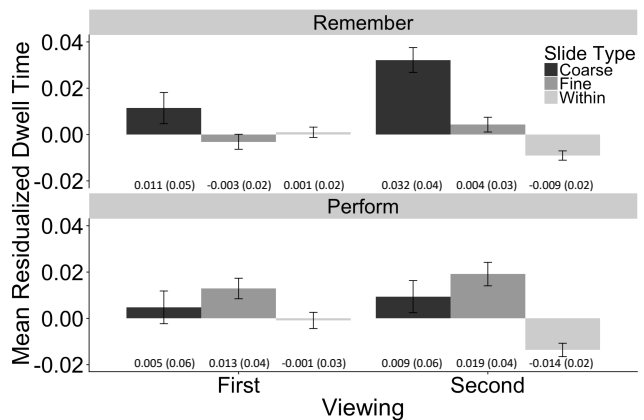


Figure 2. Average residualized dwell times (+/- SE) to coarse, fine, and within-unit slides across instructions and viewing. Means and SDs reported below bars.

Discussion

Offering participants a processing goal (i.e., to remember or perform) elevated their overall attention to a novel activity sequence relative to previous studies in which participants viewed the same event in the absence of a processing goal. Interestingly, in some respects attentional profiles were unaffected by processing instructions. For example, we found that observers targeted boundary slides with increased attention, and that this effect was robust across instructions to remember or perform. In other respects, different instructions yielded unique attentional profiles. For example, when given *remember* instructions, observers increased attention to boundaries at the coarse-grained level and did not preferentially attend to distinctive or non-distinctive regions. Conversely, when participants were given *perform* instructions, they targeted both distinctive regions and fine-level event boundaries with increased attention.

In the current study, coarse-level boundaries occurred only in non-distinctive regions of activity. Therefore, it was not possible to directly disentangle whether participants given the *perform* instructions were specifically targeting fine- over coarse-level event boundaries or if the higher dwell times to fine-level boundaries were simply elicited by observers'

increased attention to the distinctive region. Though challenging in naturalistic action, future work would benefit from using a variety of activity sequences that contain both coarse- and fine-level boundaries across distinctive and non-distinctive regions of novel activity sequences.

Despite the above-mentioned limitation, it can be concluded that a simple one-word difference in instructions elicits changes in online processing of a novel activity. Why might this occur? Recently, Flores, Bailey, Eisenberg, and Zacks (2017) demonstrated that simply instructing observers to segment activity with respect to boundaries resulted in improved memory for event sequences. Perhaps instructions to *perform* function similarly, eliciting increased attention to fine-level event boundaries (relative to instructions to *remember*) and perhaps, consequently, improvements in memory and performance.

An important next step is thus an exploration of the ways in which instructions to *remember* or *perform* and the resulting attentional patterns influence observers' ability to learn the novel shoelace tying activity. Additionally, if differences in instructions influence the granularity at which observers attend to novel action, it seems probable that memory would reflect such differences in processing. For example, those instructed to *remember* might be more likely to recall gist-level details about an activity, but little about more fine-grained information, while those instructed to *perform* might exhibit the opposite pattern and perhaps be more likely to recall fine-level details. We are currently investigating these possibilities.

The findings we report here are the first steps to understanding how differences in context influence observers' attention to unfolding activity. We additionally showcase the value of the dwell time paradigm for gathering detailed information about the consequence of such attentional influences. These results set the stage for asking a variety of new questions about how learners acquire facility with novel actions.

Acknowledgements

The authors would like to thank Rose Maier, Jenny Mendoza, Robbie Ross, Shahar Shirtz, and Jason Wallin for valuable conceptual input regarding this research.

References

- Bailey, H. R., Kurby, C. A., Giovannetti, T., & Zacks, J. M. (2013). Action perception predicts action performance. *Neuropsychologia*, *51*(11), 2294-2304.
- Blakemore, S. J., & Decety, J. (2001). From the perception of action to the understanding of intention. *Nature Reviews Neuroscience*, *2*(8), 561.
- Flores, S., Bailey, H. R., Eisenberg, M. L., & Zacks, J. M. (2017). Event segmentation improves event memory up to one month later. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(8), 1183.
- Hard, B. M., Recchia, G., & Tversky, B. (2011). The shape of action. *Journal of experimental psychology: General*, *140*(4), 586.
- Kosie, J.E. & Baldwin, D.A. (2016). A twist on event processing: Reorganizing attention to cope with novelty in dynamic activity sequences. *Proceedings of the 37th annual meeting of the Cognitive Science Society*, Philadelphia.
- Kosie, J.E. & Baldwin, D.A. (under revision). Attention rapidly reorganizes to structure in a novel activity sequence. Unpublished manuscript, University of Oregon.
- Kurby, C. A., & Zacks, J. M. (2008). Segmentation in the perception and memory of events. *Trends in cognitive sciences*, *12*(2), 72-79.
- Kurby, C. A., & Zacks, J. M. (2011). Age differences in the perception of hierarchical structure in events. *Memory & cognition*, *39*(1), 75-91.
- Newton, D. (1973). Attribution and the unit of perception of ongoing behavior. *Journal of Personality and Social Psychology*, *28*(1), 28.
- Newton, D., & Engquist, G. (1976). The perceptual organization of ongoing behavior. *Journal of Experimental Social Psychology*, *12*(5), 436-450.
- Pearce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of neuroscience methods*, *162*(1), 8-13.
- Ross, R. A., & Baldwin, D. A. (2015). Event Processing as an Executive Enterprise. *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource*.
- Sargent, J. Q., Zacks, J. M., Hambrick, D. Z., Zacks, R. T., Kurby, C. A., Bailey, H. R., ... & Beck, T. M. (2013). Event segmentation ability uniquely predicts event memory. *Cognition*, *129*(2), 241-255.
- Schwan, S., & Garsoffky, B. (2004). The cognitive representation of filmic event summaries. *Applied Cognitive Psychology*, *18*(1), 37-55.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, *28*(9), 1059-1074.
- Tanaka, Y., & Baldwin, D. (in preparation) Implicit measures of event segmentation using pupillary response. Unpublished manuscript, University of Oregon.
- Zacks, J. M., Braver, T. S., Sheridan, M. A., Donaldson, D. I., Snyder, A. Z., Ollinger, J. M., ... & Raichle, M. E. (2001). Human brain activity time-locked to perceptual event boundaries. *Nature neuroscience*, *4*(6), 651-655.
- Zacks, J. M., Tversky, B., & Iyer, G. (2001). Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General*, *130*(1), 29.
- Zacks, J. M., Speer, N. K., Vettel, J. M., & Jacoby, L. L. (2006). Event understanding and memory in healthy aging and dementia of the Alzheimer type. *Psychology and aging*, *21*(3), 466.
- Zacks, J. M., & Swallow, K. M. (2007). Event segmentation. *Current directions in psychological science*, *16*(2), 80-84.
- Zacks, J. M., Kurby, C. A., Eisenberg, M. L., & Haroutunian, N. (2011). Prediction error associated with the perceptual segmentation of naturalistic events. *Journal of Cognitive Neuroscience*, *23*(12), 4057-4066.