

The End's in Plain Sight: Implicit Association of Visual and Conceptual Boundedness

Jonathan Wehry (jonwehry14@gmail.com)

Department of Psychology, University of California, Berkeley
2121 Berkeley Way, Berkeley, CA 94720

Alon Hafri (ahafri@gmail.com)

Departments of Cognitive Science and Psychological & Brain Sciences, Johns Hopkins University
237 Krieger Hall, 3400 N. Charles Street, Baltimore, MD 21218

John Trueswell (trueswel@psych.upenn.edu)

Department of Psychology, University of Pennsylvania
425 S University Ave, Philadelphia, PA 19104

Abstract

What are the categorical distinctions shared between conceptual and visual representations? One distinction may be between bounded and unbounded entities. Previous research in sign language has shown that even non-signers associate signs with repetitive motion with atelic verbs, such as “run”, and signs with sudden motion with telic verbs, such as “arrive”. In our first study, we show this distinction holds even when the visual stimuli depicted bear no intrinsic linguistic reference: we used non-linguistic random dot motions. In our second study, we demonstrate this association occurs spontaneously, even when subjects are not making explicit semantic judgments about verbs. We use a cross-modal lexical decision task in which verbs and non-words appear superimposed on bounded or unbounded dot stimuli. We find congruency when the motion boundedness matches the conceptual boundedness of the verb. Together, these studies provide evidence for an automatic link between visual and conceptual boundedness in the mind.

Keywords: telicity; motion perception; visual boundedness

Introduction

Language allows us to describe our own perceptual experience and understand the experiences of others. As social creatures, this is critical. Thus, an understanding of how and what information is common between language and perception is both interesting and important. One category of information that might be shared across these two systems is boundedness. In the current work, we investigated whether the visual perceptual system makes a distinction between bounded and unbounded stimuli, and whether this distinction is common across visual and linguistic experience (i.e., through verbs and verb phrases).

To understand what we mean by visual boundedness, imagine you're observing an event (e.g., something moving across your field of vision), but everything is blurred so that you cannot make out objects or what category of event is taking place. Because the low-level motion properties of the scene are preserved, you could still perceive the motion properties of the event, such as whether it started and stopped. In other words, there are perceptual correlates of

boundedness even when you do not have access to high-level information about objects, goals, or events.

A second form of boundedness is telicity, or conceptual boundedness. Telicity is a similar concept to visual boundedness but in the linguistic domain. Telicity refers to whether an event described by a verb or verb phrase is construed as having an intrinsic endpoint (telic) or an undefined one (atelic; Vendler, 1957). For instance, “run” is an atelic verb. While a person could not run forever, the verb itself does not entail an endpoint. This is as opposed to a verb such as “arrive.” There is a definite endpoint entailed by the verb such that the event has only occurred when someone arrives at their destination. A simple test for this distinction is to probe the felicity or grammaticality of a sentence when adding the phrases “for an hour” (atelic) and “in/within an hour” (telic; Todorova, Straub, Badecker, & Frank, 2000). For example, one could say someone *ran for an hour* but could not say someone *ran in an hour*; conversely, it is infelicitous to say someone *arrived for an hour* but fine to say someone *arrived within an hour*.¹

One way to investigate the link between visual and conceptual boundedness is through sign language, as sign language is inherently both linguistic and visual. In Malaia & Wilbur (2012), signers were instructed to produce signs for verbs, and motion capture technology was used to record the maximum deceleration, maximum velocity, and duration of the signs. It was found that the motion properties of the signs for atelic verbs, e.g. “run”, are consistent with one another, and visually distinct from telic verbs, e.g. “arrive”. Such findings suggest that signs carry information about verb telicity iconically in the form of the sign itself.

However, although this study showed that a difference in telicity may be visually distinct, it does not indicate whether humans have access to this boundedness distinction (whether implicitly or explicitly). Strickland et al. (2015) addressed this issue: they demonstrated that even among non-signers, there is an implicit bias to map atelic signs (i.e.

¹ However, this rule is not absolute. For example, the telic verb *die* can be used with both “in an hour”, an instantaneous event, and “for an hour”, an extended process with an undefined endpoint.

signs for atelic verbs) onto atelic verbs and telic signs (i.e. signs for telic verbs) onto telic verbs. In that study, English-speaking individuals without sign language experience were shown an atelic or a telic sign and were forced to choose one of two verbs that they believed matched the meaning of that sign. For example, participants viewed the sign for “float” (an atelic verb) and were asked to choose between two words that differed in telicity (e.g., “float” vs. “leave”). Participants significantly preferred the verb that matched the sign in telicity, even when neither verb referred to the true meaning of the sign (e.g., “talk” vs. “buy”), and even for verbs with no visual correlate (e.g., “think” vs. “decide”). This shows that the human mind has access to boundedness information in visual input and can associate it implicitly with word meanings that are conceptually bounded, even though iconicity for telicity does not exist in their own language.

Although the Strickland et al. (2015) results are compelling, the scope of their conclusions is limited to perception within linguistic communication. Sign language is inherently linguistic and referential. Thus, participants can presume that these visual cues have specific linguistic meanings. This raises the question of whether these results only hold when people are performing a task where they must map from one language to another (even if one of the languages is a visual one that they have no knowledge of).

We address this question here via a new set of experiments. We used visual stimuli that were not overtly referential. Participants were shown non-linguistic motion composed of scrambled dots (extracted from biological motion stimuli) that could not be recognized as interpretable events, but nevertheless contained motion information consistent with bounded or unbounded events.

In the first experiment, participants were asked to make atelic vs. telic verb choices after viewing the visual stimulus, just as in Strickland et al. (2015). As these random dot motions are not linguistic or referential, positive findings would offer strong support for a connection between visual and conceptual boundedness. In the second experiment, we test whether such an association is automatic. In a cross-modal lexical decision task, we observed a congruency effect, such that participants were faster to confirm that a stimulus was a word when the background motion matched the boundedness of the displayed verb.

Experiment 1: Verb-Motion Matching

In this experiment, each trial consisted of the participant viewing a 3-second video clip of scrambled moving dots derived from biological motion stimuli, after which the participant had to indicate which of two visually presented verbs (one atelic and one telic) best described the clip. The video clip was designed to depict an unbounded or a bounded event as determined by separate ratings of the repetitiveness of the motion. It was predicted that participants would be more likely to select telic verbs for bounded events and atelic verbs for unbounded events.

Following Strickland et al. (2015), effects of motion boundedness were tested within three different semantic domains of verbs: Physical (e.g., *fly* vs. *hit*), Social (e.g., *argue* vs. *give*) and Mental Verbs (e.g., *think* vs. *decide*). If effects hold for all three types of verbs it suggests that motion boundedness is linked to the abstract notion of telicity rather than, for example, spatial-motion aspects of the events denoted by these verbs.

Method

Participants Twenty-four participants were recruited from the University of Pennsylvania undergraduate body and participated for course credit. This was the same number as in the Strickland et al. (2015) study as their effect size was not available for a power analysis. All participants were fluent speakers of English with normal or corrected to normal vision.

Visual Materials A personal computer running the Psychophysics Toolbox Version 3 for MATLAB was used to run this experiment (Kleiner et al., 2007). Sixty biological motion, or biomotion, videos of three seconds from the CMU Graphics Lab Motion Capture Database were used with the BioMotion Toolbox (van Boxtel & Lu, 2013).

Biomotion videos are produced via motion capture, whereby each joint on a person’s body is attached to a sensor. The positions of these sensors are then recorded during movement. This produces a video, composed of dots, in which the overall shape, size, and movement of an individual is maintained but the fine details and body form are removed.

Crucially, in our versions, body structure information was removed from these videos while preserving the overall motion signal. This was done by randomizing the start point of each individual dot, but then preserving its relative motion path from that start point. For example, the dot that corresponded to the person’s right elbow may, at the start of the animation, be located to the left of where their left ankle was and the dot corresponding to their left ankle may now start right above where their right knee was. This removes the benefits of being able to tell what action is occurring (because the intact structure of the body is removed) and ensures that participants only get information about the motion properties of the dots, e.g., velocity and acceleration.

Selection and Norming of Video Materials Videos were initially selected by JW, and then their boundedness was confirmed using a norming procedure. JW rated a random set of 574 scrambled videos from the CMU database on perceived boundedness. Subsequently, we presented 79 unbounded and 61 bounded candidate videos to twenty-two undergraduates in a norming study. Although the CMU database includes descriptions for each video, they were ignored for video selection. Using these ratings, we chose a set of 60 videos to use in the subsequent experiments.

Participants were asked to rate each video for repetitiveness or deceleration (between subjects) on a scale

of 1 to 7, e.g. “Rate the video based upon how repetitive you think the motion is” or “based upon how fast you think the motion decelerates.” These properties were used as they were found to be indicators of boundedness in previous studies (Malaia & Wilbur, 2012; Strickland et al., 2015).

Although our intention was to define bounded videos as those with the highest deceleration and least repetition in the motion, deceleration ratings proved inconclusive as across all videos there was little deviation from the average of 4. That is, across all 140 videos participants tended to choose a middle value ($SD = 0.85$). This was perhaps due to the difficulty of the task and the nature of the stimuli (independently moving dots). In contrast, repetition ratings had high variety across items and participants ($SD = 1.51$).

As a result, only the repetition ratings were used to select the sixty videos for the main study. The sixty videos were selected by taking the 40 videos with the highest average ratings and the 40 videos with the lowest average ratings and then sorting these videos by lowest standard deviation across ratings. The 30 videos from each group with the lowest standard deviation were then selected. We considered the videos with high repetition ratings as Unbounded and videos with low repetition ratings as Bounded. The mean repetition rating for unbounded videos was 5.70 and for bounded videos was 1.91, and the two groups differed reliably ($p < 0.0001$).

Verbal Materials Five atelic and five telic verbs were chosen from each of three separate conceptual domains: physical, social, and mental. This resulted in fifteen telic and fifteen atelic verb pairs. Each pair consisted of one telic and one atelic verb, approximately matched for log frequency. Fourteen of the 18 Strickland et al. verbs were used (4 not used due to low frequency), with an additional sixteen verbs (seven telic and nine atelic) generated by author JW to maintain approximate match in log frequency. The verbs were the following. Atelic: *run, fly, play, paint, sing, think, consider, imagine, dream, study, talk, discuss, fight, love, argue*; Telic: *enter, die, leave, hit, grab, decide, accept, forget, choose, remember, marry, sell, buy, give, and take*. To create each participant’s list of paired verbs for each trial, the atelic and telic verbs for each domain were shuffled and then paired (within domain) to produce fifteen total pairs (five for each domain). This shuffling was performed four times (60 trials in total). Verb pairs (trials) were shuffled and paired randomly with videos.

Procedure During the experiment, participants were instructed that they would be shown a short clip of moving dots and asked to choose which of two verbs better fit the clip. They were told to use their intuition to make their verb choice and that there was no right answer. After the instructions, two practice trials were given.² On each trial, participants were shown the video twice before making their selection, to ensure they could adequately perceive the

² Thus, there were 2 practice videos and 58 trial videos. Later “test” trials using these videos were discarded from analysis.

motion. The video slowly faded in over the first half second and then faded out during the last half second to diminish influences of motion onset. After the video disappeared, the two verbs appeared on the screen, one on each side. Participants then made their selection (f for left or j for right). See Figure 1 for a schematic of trial types.

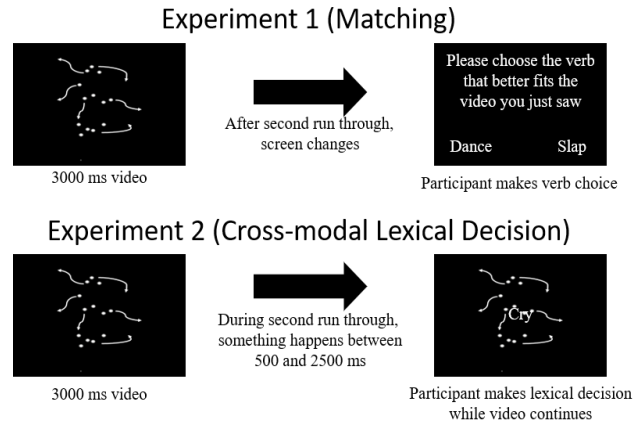


Figure 1: Trial schematics for Experiment 1 (explicit matching) and 2 (lexical decision). In Exp. 1, participants watched a 3 second video twice, then made a choice of which of two verbs better fit the video. In Exp. 2, during the second viewing, they were instead given a lexical decision (word/nonword) or attention task. Lines with arrows illustrate motion paths and were not seen by subjects.

Results

Figure 2 presents the mean proportion of telic verbs selected (subject means) as a function of whether the motion event was unbounded or bounded, overall and separately for each verb domain. As predicted, telic verbs were selected less often for Unbounded motion events ($M=0.39, SE=0.02$) than for Bounded motion events ($M=0.55, SE=0.02$). Moreover, the effect of motion Boundedness was very similar for each verb domain.

To test for reliability, we used a multilevel logistic regression to model binary trial-level choices for the telic verb. Fixed effects consisted of motion Boundedness (bounded vs. unbounded), Verb type (physical, social or mental), and the interaction. The maximal random effects structure was used for each subject and a random intercept was used for each item (each video). The significance of factors was performed by comparing likelihood-ratio values for nested models that included main effects and interactions of factors to models without them.

The best-fitting model showed a reliable effect of motion Boundedness ($\beta=0.339, SE=0.096, z=3.531, p=0.0004$), but no reliable effects or interactions with Verb type. Removing Verb type and the interaction from the model did not decrease the fit of the model ($\chi^2(4)=4.38, p=0.346$), but further removing the effect of Boundedness did ($\chi^2(1)=27.92, p<0.0001$) – indicating that the motion Boundedness of the videos reliably predicted telic choices.

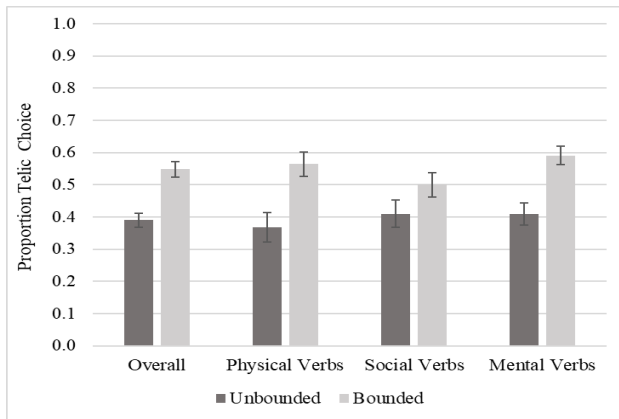


Figure 2: Mean proportion of telic verb choices as a function of Motion Boundedness (Unbounded, Bounded), overall, and by verb domain. Average of Subject Means. Error bars indicate ± 1 Standard Error.

One concern is that people may have developed a strategy over time or discovered the purpose of the experiment and acted accordingly. Although unlikely, as no participant revealed an explicit awareness of the hypothesis or purpose when questioned, we examined the time course of the effect. If it was an explicit strategy, we might expect a difference between the beginning and the end of the experiment. We tested whether there was a statistical difference between the first twenty trials and last twenty trials regarding whether participants showed a differing proportion of telic verb preference for bounded and unbounded videos. This test showed no significance (paired $t(23) = -0.635, p = 0.532$), suggesting a strategy did not emerge during the experiment.

Discussion

The current results show that the visual system can distinguish between “bounded” and “unbounded” motion stimuli, as defined by repetitiveness. Furthermore, people are biased to associate visual boundedness with conceptual boundedness in verbs. Surprisingly, this occurs even when these verbs have no visual manifestation, as is the case for the domain of mental verbs. This implies that people were implicitly encoding boundedness and telicity when observing these biomotion stimuli. Put another way, these results show that the distinction of boundedness that is present in both the visual and linguistic systems is shared or otherwise accessible by the two systems. This extends Strickland et al. (2015) by demonstrating that this association is not just due to the referential task (associating two linguistic items), but rather it exists even when using visual stimuli that have no inherent referential properties (i.e. scrambled biomotion stimuli). Thus, it appears that conceptual boundedness is a basic property of the visual and linguistic systems.

Although we find these results to be compelling, it would be of interest to understand if similar results can be obtained when participants are not attempting to link the video of the motion event to the linguistic stimuli. That is, could such a

connection between motion boundedness and linguistic telicity arise spontaneously, even when linguistic judgements do not involve connecting the verb to the motion video? If so, it would suggest that the perception of motion boundedness automatically activates linguistic telicity. We explore this issue in Experiment 2.

Experiment 2: Cross-Modal Lexical Decision

In this experiment, we present preliminary results from two versions of a cross-modal lexical decision task. Each trial involved the participant viewing the same clips used in Exp. 1. However, participants were not presented with a forced choice task. Instead, for target trials, a single telic or atelic verb appeared centrally over the video, and the participant’s task was to make a lexical decision (word / nonword). The core prediction of this experiment is the following: *If the perception of motion boundedness spontaneously activates linguistic telicity, then we would expect a congruency effect, such that Bounded motion events would speed judgments of telic verbs whereas Unbounded motion events would speed judgments of atelic verbs.*

The results are preliminary because each version of the experiment suffered from some issues related to response time measurement precision, such that we may have been underpowered to detect robust results in either alone. Nevertheless, across versions, the patterns are significant and consistent with our hypothesis. Thus, we present both sets of results together, noting differences between versions below as needed. Further versions of the experiment with more precise RT measurements are planned.

Differences between Versions of the Experiment Each version of the experiment suffered from precision issues with RT collection. In version 1 of the experiment, participants used a keyboard to make their responses (f for word, j for nonword, or spacebar if the dots changed color). However, due to a coding error, responses were only recorded at each screen refresh (every 16.67 ms); thus, these measurements were imprecise. In version 2, we used an E-prime button box instead of keyboard, since it is known for its measurement precision. Mean accuracy and mean RTs were nearly on par across experiments. However, the buttons on the button box used in version 2 were differentially sensitive; 4-5% of trials were timeouts, while those in version 1 were nearly 0%, suggesting that sometimes the buttons were not responsive.

Other differences were the following. In version 1, unbounded videos were randomly assigned one of the onset times for the bounded video. In version 2, this assignment was fixed for each list. That is, in version 2, unbounded video A would always have onset time of bounded video 1, video B would have onset time of bounded video 2, etc.

Method

Participants Experiment 2 consisted of two separate versions. A total of 116 subjects were recruited from a

university's undergraduate student body. Participants were excluded based on pre-determined criteria: accuracy below 80% on any of the three tasks (described below). After exclusion, there were 54 in the first version and 47 in the second. For both versions, the goal sample size was determined by doubling the sample size of experiment 1. The data from participants 49-54 from version 1 were included due to this experiment being underpowered (described below). Students signed up to participate in exchange for study credit. All participants were fluent English speakers with normal or corrected to normal vision.

Materials The same materials from Experiment 1 were used (i.e., the same fifteen telic and atelic verbs across three domains and the same sixty test videos).

An additional sixty videos from the CMU Graphics Lab database were used for filler attention trials (described below). An additional twelve videos and four verbs were used for practice trials. Sixty-four non-words were created using Wuggy, a word generation tool that creates nonwords matched with inputted real words on phonotactics and word length (Keuleers & Brysbaert, 2010). These non-words were used for the lexical decision task (described below).

Although our videos were rated for overall boundedness in the previous norming study, this does not indicate *when* in the video a boundary occurred. To ensure that the onset time of the word stimulus coincide with a motion boundary (for bounded videos), study authors JW and AH and research assistants chose the boundary point by watching each video frame by frame. In version 1, JW and AH individually watched and chose the points. If the differences between the two values was more than thirty frames, JW re-watched the video and made the final decision. If the difference was less than thirty frames, the average of the two was taken. In version 2, to get a more reliable value for boundedness point, median frame values for each video across an additional four research assistants, in addition to JW and AH, were taken. The average difference between the first and second experiment version was three frames. Since each video was only 3 seconds long, boundaries closer than 0.5 seconds towards the beginning or end of the video were constrained to be at 0.5 or 2.5 seconds, respectively. Since they do not have a motion boundary, the distribution of onset times for unbounded videos and filler (attention) videos were matched to the bounded videos.

Procedure During the experiment, there were three trial types: two were lexical decision (word or non-word). In these trials, participants simply had to press one button if the string of letters that appeared was a word, and another if it was not. The third trial type was an attention catch task, designed to ensure participants attended to the visual dot stimulus. In this catch task, dots would briefly change color from white to blue (0.5 sec); thus, this task made no reference to visual motion.

Participants were instructed to press a different key for each of these three trial types. On each trial, one of the three visual changes would appear: a word superimposed on the

dot stimuli, a non-word on the dot stimuli, or the color change type. Participants were not made aware of what the current or next trial would be ahead of time. After the instructions were given, twelve practice trials (four of each type) were given. On each trial, participants were shown the 3-second video twice to ensure they could adequately perceive the full motion. The visual stimulus appeared at the pre-determined onset time during the second viewing. The videos faded in and out over the first and last half second to avoid sudden visual transients.

The visual stimuli (word, non-word, color change) appeared either *at* the pre-selected onset frame, or 0.25 sec *after* (counterbalanced with each condition). However, this timing factor was collapsed over in analyses.

Each verb and nonword was paired once with a bounded video and once with an unbounded video. There were 192 total trials (60 word, 60 non-word, 60 catch, 12 practice).

Results

The results for the first and second version of the experiment are being presented together. For an explanation and discussion of this decision, see below. Reaction times ± 2.5 SDs from each subject mean were excluded (2.9% of trials), as well as timeout trials. Word trials: Accuracy 94.8% (*SD* 3.5%), mean RTs 676 ms (*SD* 119ms); Non-word trials: Accuracy 92.3% (*SD* 5.1%), mean RTs 808ms (*SD* 155ms); Attention trials: Accuracy 94.1% (*SD* 4.0%), mean RTs 629ms (*SD* 139 ms). All statistical analyses were performed on inverse RTs ($-1000/RT$), on accurate word trials only. Inverse RTs were used to improve normality of RT distributions for model fitting (Baayen & Milin, 2010).

Figure 3 presents the inverse reaction time (subject means) as a function both of verb telicity (telic vs. atelic), and motion boundedness (whether the visual stimulus was unbounded or bounded). As predicted, we observe an interaction between verb telicity and visual boundedness on reaction times for lexical decision: reaction times to atelic verbs were faster when the visual motion was unbounded, and reaction times to telic verbs were faster when the visual motion was bounded. However, the effect here was subtler than in Exp. 1, as should be expected: participants were not performing an explicit matching task but deciding whether the word that appeared was a real word or a non-word (or were performing a color change detection task, for filler trials).

To test for reliability, we used a multilevel linear regression to model inverse reaction time. Fixed effects consisted of Verb Telicity (telic vs. atelic), Motion Boundedness (bounded vs. unbounded), Verb Type (physical, social or mental), and all relevant interactions. To account for mean RT differences in experiment versions, a main effect of experiment Version was also included. The maximal random effects structure that converged was used for each subject (random intercept and random slopes for telicity and motion boundedness), and a random intercept was used for each verb. We compared nested models with and without these factors and interactions.

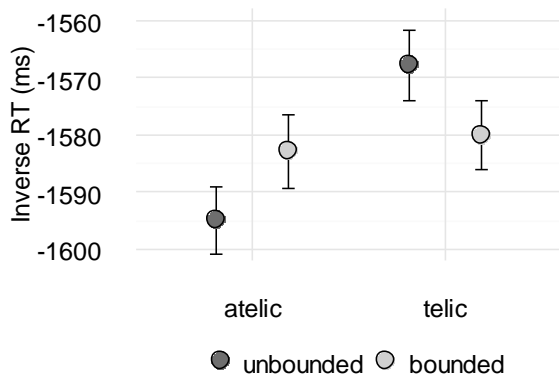


Figure 3: Mean inverse response times (-1000/RT) on lexical decision trials (word only), as a function of verb telicity (telic, atelic) and motion boundedness (unbounded, bounded). Average of Subject Means. Error bars indicate ± 1 Standard Error. Up indicates slower RTs, down faster RTs, as is standard for raw RT plots.

The best fitting model was one that included main effects and interactions of Verb Telicity and Motion Boundedness, and Verb Type and Motion Boundedness. This model produced a reliable interaction between Verb Telicity and Motion Boundedness ($\beta=6.56$, $SE= 2.91$, $t(5555)=2.26$, $p=0.024$). Telicity and Motion Boundedness were contrast coded in the following way: telic=1, atelic=-1; unbounded=1, bounded=-1. Thus, the positive β indicates greater RTs for the mismatched conditions (e.g. telic+unbounded or atelic+bounded). Adding a triple interaction of Telicity, Boundedness, and Verb Type only marginally improved the fit ($\chi^2(4)=8.34$, $p=0.08$). Further, removing the interaction of Telicity and Boundedness *did* decrease the fit of the model ($\chi^2(1)=5.09$, $p=0.024$). This confirms the interaction of Verb Telicity and Motion Boundedness that we expected.

Discussion

Experiment 2 results produced a pattern of congruency that would be expected if the perception of motion boundedness activated linguistic telicity. Reaction times to atelic verbs were faster when preceded by Unbounded as opposed to Bounded motion events whereas the opposite was found for telic verbs. That this effect arose even though participants were not trying to relate the verbs to the motion videos suggests that the relation between boundedness and telicity occurs spontaneously, without conscious effort.

There are a few important issues that force us to see these results as preliminary and necessary of replication. Even though the timing issues mentioned in the Method section merely decrease RT precision, it is also the case that the results of both versions were pooled post-hoc, after discovering these RT precision issues. Although these are not small issues that should be overlooked, we believed it was still worthwhile to report on these data as they are

promising, consistent with the judgment data of Experiment 1, and spur the need for replication.

Although we did not find strong evidence for differences among verb domains in the congruency effect, future replications with higher power will allow us to determine whether this difference is real, and if so, whether only certain of the domains (e.g. physical) demonstrate the effect. Nevertheless, results of Experiment 1 (matching) and Experiment 2 (lexical decision) do suggest that the congruency effect is general, regardless of domain.

General Discussion

We have shown here that the motion properties of what can only be characterized as scrambled moving points of light yield systematic and expected interpretive responses from observers, concerning the detection of motion boundedness, and its relation to conceptual and linguistic telicity.

In Experiment 1, we observed that participants were more likely to choose a telic verb over an atelic verb to describe a bounded non-repetitive motion, even when the meanings of these verb pairs denoted abstract mental events (e.g., *think* vs. *decide*). This extends the Strickland et al. (2015) findings to visual stimuli more generally, even stimuli without linguistic and referential properties (signs). Experiment 2 offered preliminary results that activation of conceptual telicity from motion signals arises automatically, such that telic verbs show a congruency effect with bounded motion and atelic verbs show an effect with unbounded motion. Activation of concepts from motion signals has been observed before in similar tasks, e.g., that upward motion will speed reaction time to *rise* as opposed to *fall* (Meteyard et al., 2008). Additionally, previous research has shown that the ability to judge an action verb on a lexical decision task is correlated with the ability to judge a non-scrambled point-light action on an action decision task, e.g. is this a valid human action (Bidet-Ildei & Toussaint, 2015). What is surprising in the present study is that activation is for a highly abstract categorical feature (boundedness) that arises from seemingly continuous motion signals, and that the activation affects judgments even for verbs labeling events in the social and mental domains, which have no overt visual boundedness cues.

Observing an implicit association between visual and linguistic boundedness suggests there is an underlying amodal conceptual distinction that both systems have access to: a distinct categorical representation of boundedness, which may indeed be a conceptual primitive similar to that proposed for causation (see, e.g., Jackendoff, 1996; Rolfs, Dambacher, & Cavanagh, 2013). Of course, a representation of boundedness is not limited to the event domain; things may be conceived of as objects or substances, which have perceptual consequences of their own (vanMarle & Scholl, 2003). In future work we plan to refine our experimental tasks to replicate our lexical decision effects in the event domain, and to determine if they extend to boundedness across conceptual domains (events and objects).

References

- Baayen, R. H., & Milin, P. (2010). Analyzing Reaction Times. *Int. Journal of Psychology Research*, 3, 12–28.
- Bidet-Ildei, Christel & Toussaint, Lucette. (2014). Are judgments for action verbs and point-light human actions equivalent?. *Cognitive processing*. 10.1007/s10339-014-0634-0.
- Jackendoff, R. (1996). Conceptual semantics and cognitive linguistics. *Cognitive Linguistics*, 7(1), 93-129.
- Keuleers, E., & Brysbaert, M. (2010). Wuggy: A multilingual pseudoword generator. *Behavior Research Methods* 42(3), 627-633.
- Kleiner M, Brainard D, Pelli D, 2007, "What's new in Psychtoolbox-3?" *Perception* 36 *ECVP*.
- Malaia, E., & Wilbur, R. B. (2012a). Kinematic signatures of telic and atelic events in ASL predicates. *Language and Speech*, 55(3), 407–421.
- Meteyard, L., Zokaei, N., Bahrami, B., & Vigliocco, G. (2008). Visual motion interferes with lexical decision on motion words. *Current Biology*, 18(17).
- Rolf, M., Dambacher, M., & Cavanagh, P. (2013). Visual adaptation of the perception of causality. *Current Biology*, 23(3), 250-254.
- Strickland, B., Geraci, C., Chemla, E., Schlenker, P., Kelepir, M., & Pfau, R. (2015). Event representations constrain the structure of language: Sign language as a window into universally accessible linguistic biases. *Proceedings of the National Academy of Sciences*, 112(19), 5968–5973.
- Todorova, M., Straub, K., Badecker, W., & Frank, R. (2000). Aspectual coercion and the online computation of sentential aspect. In *Proceedings of the Cognitive Science Society* (Vol. 22).
- van Boxtel, J. J. A., & Lu, H. (2013). A biological motion toolbox for reading, displaying, and manipulating motion capture data in research settings. *J. of Vision*, 13(12), 7–7.
- vanMarle, K., & Scholl, B. J. (2003). Attentive Tracking of Objects Versus Substances. *Psychological Science*, 14(5), 498–504.
- Vendler, Z. (1957). Verbs and Times. *The Philosophical Review*, 66(2), 143–160.