



Linguistic consequences of event segmentation in visual narratives: implications for prominence

Elsi Kaiser

To cite this article: Elsi Kaiser (2019): Linguistic consequences of event segmentation in visual narratives: implications for prominence, Language, Cognition and Neuroscience, DOI: [10.1080/23273798.2019.1667000](https://doi.org/10.1080/23273798.2019.1667000)

To link to this article: <https://doi.org/10.1080/23273798.2019.1667000>



Published online: 22 Sep 2019.



Submit your article to this journal [↗](#)



Article views: 50



View related articles [↗](#)



View Crossmark data [↗](#)

Linguistic consequences of event segmentation in visual narratives: implications for prominence

Elsi Kaiser

Department of Linguistics, University of Southern California, Los Angeles, CA, USA

ABSTRACT

This experiment uses comic-type visual narratives to investigate effects of event segmentation on the representation of causal events (with an agent and a patient) in discourse. Prior work leaves open the question of whether the prominence of entities and events in mental models of discourse can be dissociated. By presenting participants with one-panel vs. two-panel segmentations of the same event, we tested whether differences in event segmentation boost the prominence of the patient, the consequence event, or both, as reflected in people's choices about what is mentioned next. The results show participants are more likely to mention consequences in two-panel than one-panel conditions, indicating that panel segmentation influences the cognitive prominence of event-level representations. However, we find no clear evidence for segmentation influencing the cognitive prominence of entities. This suggests language users separately track expectations about who is mentioned next and expectations about what kind of event is described next.

ARTICLE HISTORY

Received 28 August 2018
Accepted 21 August 2019

KEYWORDS

Discourse structure; event segmentation; visual narrative; prominence; causality

Introduction

The world consists of a continuous stream of activity, which we segment into meaningful events. Event segmentation involves the recognition of event boundaries and relations between events (e.g. Radvansky & Zacks, 2017; Zacks & Swallow, 2007; Zacks, Speer, Swallow, Braver, & Reynolds, 2007). Understanding events and the relations between them is also crucial for discourse-level linguistic processing (e.g. Hobbs, 1979; Kehler, Kertz, Rohde, & Elman, 2008; Lascarides & Asher, 1993; Mann & Thompson, 1988), but the input is different: Language consists of discrete, compositional elements in contrast to the on-going influx of information in the world. When a speaker describes a sequence of events, she is faced with linguistic choices about how to describe each event and the entities involved. The present paper investigates causal sequences to gain insights into how visual cues to event segmentation influence linguistic choices about what to mention next, and how this can inform our understanding of the cognitive prominence of entities and events.

Consider a causal sequence where Lisa sticks out her foot (the *cause*) and this makes Mary trip (the *consequence*). Causality is fundamental to conceptualising the world (e.g. Pearl, 2009; Radvansky & Copeland, 2000; Wolff, 2003). The tripping event can be described

in multiple ways, e.g. with a single sentence (*Lisa tripped Mary*) or two sentences expressing the cause and consequence (e.g. *Lisa stuck out her foot. Mary fell*), see e.g. Comrie (1985), Wolff (2003).

What influences speakers' linguistic choices? One factor at play is the notion of "prominence" (see e.g. Ariel, 1990; Chafe, 1976; Givón, 1983; Gundel, Hedberg, & Zacharski, 1993) – also known as salience or accessibility. Linguists agree that some aspects of our mental representations of events and situations are more prominent than others, which has consequences for language production and comprehension (Jasinskaja, Chiriacescu, Donazzan, von Heusinger, & Hinterwimmer, 2015; see also Johnson-Laird, 1983; van Dijk & Kintsch, 1983 on mental models). In particular, it is agreed that not all *referents* are equally prominent, but less is known about the prominence of *events*. To understand mental models of discourse, it is necessary to consider both entities and the events they participate in.

In psycholinguistic work, the prominence of entities and events is often not clearly dissociated. E.g. work on implicit-causality verbs (e.g. Garvey & Caramazza, 1974; Hartshorne, 2014; Kehler et al., 2008) shows that depending on the verb, a subsequent explanation can be strongly associated with reference either to the preceding subject or object. Thus, it could be the case that the discourse processing system treats the prominence

of entities and of events as being “one and the same” – an entity is prominent when it participates in a prominent event and vice versa. This leaves open the question of whether the prominence of entities and events in mental discourse models can be dissociated – which has implications for our understanding of the notoriously complex notion of prominence and the components of discourse representation.


The present paper aims to further our understanding of these issues by investigating the prominence of events and entities in causal sequences involving agents and patients, and causal events and consequence events. The experiment makes use of the segmentation options offered by visual narrative to test whether making the causally affected entity (patient) more visually prominent affects the cognitive prominence of the *patient* entity and/or the cognitive prominence of the associated *consequence* event.

Segmenting events and entities into panels

Comics resemble language in forcing the segmentation of events into discrete segments: “To deal with the capture [...] of [...] events in the flow of the narrative, they must be broken into sequenced segments” (Eisner, 1985/2008, p. 39). How narratives are segmented into panels is a key aspect of research on comics (e.g. Cohn, 2013, 2018; Eisner, 1985/2008; McCloud, 1993; Postema, 2010; Saraceni, 2016; Stainbrook, 2016; see also Magliano, Kopp, Higgs, & Rapp, 2017).

Events can be segmented into panels in different ways. The two options relevant here are in Figure 1(a, b): The agent and patient of a causal event sequence can be represented in one panel (Figure 1(b)) or separated into two (Figure 1(a)). The blank space between the panels is called the gutter.

These two segmentation options are *two different ways of visually segmenting the same causal event*. The current experiment tests two hypotheses concerning panel segmentation, one having to do with the prominence of entities and the other with the prominence of events:



(a)

Figure 1a. Agent and patient segmented into two panels (Super Duck Comics #2, Winter 1944, <http://comicbookplus.com/?dliid=69519>).



(b)

Figure 1b. Agent and patient presented in one panel (Super Duck Comics #32, June 1950, <http://comicbookplus.com/?dliid=57706>).

Hypothesis #1: segmentation influences representation of entities' prominence

A large body of linguistic work shows that agents are privileged over patients in various respects (e.g. Dowty, 1991; Fillmore, 1968). Pioneering work on visual narratives by Cohn and Paczynski (2013) similarly found that agents are cognitively more prominent than patients. Can panel segmentation modulate this default agent prominence? Cohn (2015) describes panels as “attention units” that focus attention on the character in each panel. Thus, the two-panel variant emphasises the prominence of each character in turn. It highlights the patient by isolating it in its own panel. Thus, even if the patient is lower-prominence than the agent, the two-panel variant is predicted boost the prominence of the patient relative to the one-panel variant. A similar prediction can be derived from the Event Horizon Model (Radvansky 2012; Radvansky et al., 2011), assuming it can apply to entities. This model posits that “information that was prior to an event boundary (...) becomes less available after” the boundary (Radvansky et al 2011, p. 1632). If the gutter signals a boundary, the default agent prominence should weaken in the two panel-version.

Hypothesis #2: segmentation influences representation of consequence events' prominence

According to this hypothesis, two-panel variants boost the prominence of consequence events, relative to one-panel variants. To see why, let us first note that Cohn (2015) and Cohn and Kutas (2017) argue that two-panel variants require additional processing: Readers need to infer that the two characters in distinct panels need to be “integrated into a single spatial environment” (Cohn & Kutas, 2017, p. 5; also McCloud, 1993; Saraceni, 2016; Stainbrook, 2016). Support for the claim that two-panel variants need additional inferring comes from Cohn and Kutas (2017) ERP study showing that characters split into two frames elicit an increased positivity (P600). Other work shows that

inferencing increases processing depth and recall: Myers, Shinjo, and Duffy (1987; also Keenan, Baillet, & Brown, 1984) found that sequences of causally related sentences were recalled better when the level of causal relatedness was moderate, not high. They attribute this to highly causally related sequences requiring no additional inferencing, and thus not being processed as deeply as moderately related sequences which require inferencing (see also Olmstead & Kuhlmeier, 2015 on more elaborated processing resulting in more detailed encoding, Radvansky & Copeland, 2000 on causal connections).

Let us now return to two-panel sequences like Figure 1(a). If a reader sees one panel depicting a person yelling and another panel showing a fleeing dog, she/he needs to infer that the two characters are in the “same spatial environment” (Cohn, 2015; see also Cohn & Wittenberg, 2015; Magliano et al., 2017). This integrative inference allows the reader to understand that, in Figure 1(a), the character in the second panel is the patient affected by the causal agent’s actions – i.e. that actions have consequences that affect others (see Testelec, 1998 i.a. for linguistic work on affectedness). According to Hypothesis #2, the integrative inference about affectedness activates an abstract notion of “consequence”. Although one-panel variants like Figure 1(b) can also activate this, according to Hypothesis #2 the “consequence” notion is more prominent with the two-panel variant, due to the integrative inference (see Cohn, 2015; Cohn & Kutas, 2017; also Magliano et al., 2017) – which leads to a higher proportion of continuations that start by describing a consequence event.

Relatedly, the Event Horizon Model (Radvansky 2012; Radvansky et al., 2011) posits that causality plays a central role in event models. Thus, if the gutter renders pre-boundary information less available, it seems reasonable to predict that the abstract notion of “consequence” associated with the post-boundary affected patient becomes more available in two-panel versions. (We do not aim to adjudicate between Cohn and colleagues vs. Radvansky and colleagues. We focus on the initial empirical step of assessing how panel segmentation affects entities vs. events.)

Hypotheses #1 and #2 are not mutually exclusive: Panel segmentation could influence the cognitive prominence of *both* events and entities. However, if only events or only entities are affected, this would suggest a dissociation: the representation of prominence distinguishes between events and entities.

Production task

To measure prominence, participants wrote continuations for sentence fragments. The choice of task was

motivated by a large literature showing that topical, prominent referents are more likely to be mentioned again than non-topical, non-prominent referents (e.g. Givón, 1983, p. 15). Analyzing which referent people mention first indicates which referent is most prominent (e.g. Arnold, 2001; Grüter, Rohde, & Schafer, 2017; Stevenson et al., 1994; von Heusinger & Chiriacescu, 2013). Extending this to events means that people’s choices about which events to describe provide a measure of which event they view as prominent.

Predictions

According to *Hypothesis #1*, panel segmentation influences conceptualisation of entities’ cognitive prominence. In English, agents are typically subjects and more cognitively prominent than non-agentive non-subjects: We expect a default, overarching preference for agent-initial continuations. However, if isolating the patient in its own panel boosts its prominence, two-panel variants should elicit more patient-initial continuations than one-panel variants. According to *Hypothesis #2*, panel segmentation influences the prominence that comprehenders assign to components/properties of events. As sketched out above, the prediction is that the abstract concept of “consequence” becomes more prominent and leads to two-panel variants eliciting a higher proportion of continuations that start by describing a consequence event.

Method

Participants

Twenty-eight native speakers of American English from the University of Southern California community participated. This research was reviewed and approved by the USC Institutional Review Board.

Materials

Targets manipulated whether a transitive event with an agent and an affected patient (e.g. *kick, push, tickle, trip*) is depicted with both characters in the same panel, or split across two panels (Figure 2). The distance in pixels between the characters was held constant. Both variants have the same closing panel, which shows two new characters, one of whom says “... and then ...”. Participants wrote a continuation for this open-ended fragment, which allows continuations about the pictured event, a consequence, a non-causal subsequent event, a preceding event, and so on.

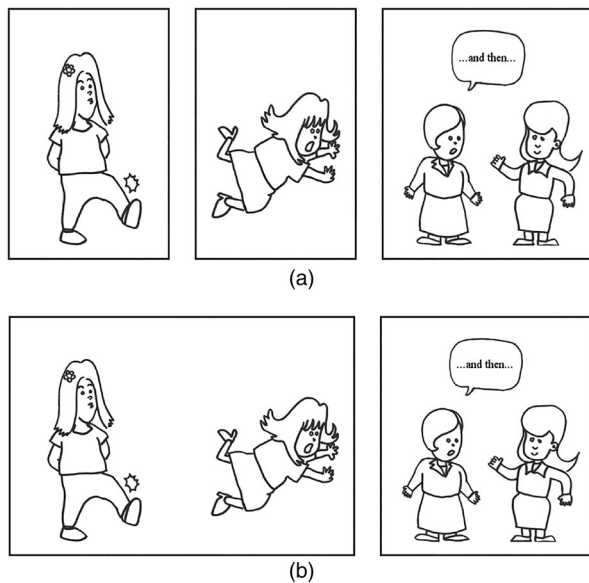


Figure 2. (a) Example target: two-panel version, (b) example target: one-panel version.

The study contained eight targets, each depicting a different verb/event, presented using a Latin-Square design. Sixteen fillers, consisting of two or three panels, were also constructed. The initial panels depicted one, two or three people. Similar to targets, fillers had a closing panel with a speech or thought bubble, with various prompts (*and then, I wonder if, etc.*).

Procedure

The sequence of two or three images was shown on the screen, and participants typed in a continuation for the fragment. Participants were told to write whatever first came to mind. Participants wrote three sentences for each item, but our focus is on the first sentence.

Data analysis

The continuations were analysed for event type, and whether the first-mentioned referent is the agent or patient of the depicted action. Data was double-coded by two native English-speaking annotators blind to experimental condition, and analysed using mixed effects logistic regression (glmer, lme4 1.1-17 in R, R Core Team 2018). Models with fully specified random effects for subjects and items were used whenever possible.

We focus on the first sentence of the three that participants wrote for each item. Story grammars research (e.g. Mandler & Johnson, 1977; Rumelhart, 1975) shows that narratives have certain basic parts, including a description of the outcome/consequence of the protagonist's

actions. Instructing participants to write only one sentence could bias them towards this outcome sentence, which could inflate the rate of consequence-type continuations. To avoid this, people wrote three sentences but only the first one was analysed.

The analyses focus on the first-mentioned entity in the continuation, following existing work (e.g. Arnold, 2001; Grüter et al., 2017; Rohde & Kehler, 2014; Stevenson et al., 1994), as that is assumed to be the cognitively most prominent referent. Participants could write active (*Bob fell over*) or passive sentences (*Bob was kicked by John*): we always analysed the first-mentioned referent.

For the **event-based analysis**, we analysed whether the sentence (i) mentioned a consequence resulting from the pictured event (*consequence-type* continuation), (ii) only described the pictured event itself (*pictured-event* continuation), or (iii) mentioned another kind of event.¹ We regard *pictured-event* continuations as the default: They simply describe what is shown in the image and do not require creation of novel content (e.g. past/future events). Examples are in Table 1. We focus on whether a consequence (e.g. dress getting ripped, an apology occurring) of the pictured event is mentioned in the first sentence, as that is most relevant for our predictions.

For the **entity-based analysis**, we coded whether the first-mentioned entity was the agent or the patient of the depicted action² (Table 1). (The first-mentioned entity was the subject of the continuation sentence 93.9% of the time; 6.1% were possessors). Continuations where the first-mentioned referent was something other than a singular third-person expression referring to the agent/patient were excluded from subsequent entity-based analyses (e.g. one of the other characters in the final panel, or a plural referent), following common practice in research on pronoun use.³ This affected 34% of the data. The number of excluded responses did not differ by condition (chi-squared test, $p > .25$).

Table 1. Examples (for “trip” item) of event type and entity type coding labels (Characters’ names were not prespecified, participants used whatever names they wanted).

Continuation	Event type	First-mentioned entity
... and then <u>Mary</u> tripped my daughter!	Pictured event	Agent
... and then <u>that girl</u> tripped Mary.	Pictured event	Agent
... and then <u>the girl</u> apologized profusely for accidentally tripping the lady.	Consequence type	Agent
... and then <u>Monica’s fancy new outfit</u> got ripped in several places.	Consequence type	Patient
... and then <u>Lucy</u> was badly scraped when she landed.	Consequence type	Patient

Results

Event-based analysis

Figure 3 shows the proportion of pictured-event continuations, consequence-type continuations and “other” continuations for the one-panel and two-panel conditions. The proportion of consequence-type continuations is higher in the two-panel than the one-panel condition ($\beta = -0.905$, $SE = 0.332$, $z = -2.724$, $p = .006$). In contrast, the proportion of pictured-event continuations is higher in the one-panel than the two-panel condition ($\beta = 0.899$, $SE = 0.34$, $z = 2.644$, $p = .008$) – suggesting a trade-off between these two continuation types. The proportion of “other” continuations does not differ between conditions ($p > .7$).

Entity-based analysis

Figure 4 shows the proportion of continuations in the one-panel and two-panel conditions where the agent of the depicted event is mentioned first vs. where the

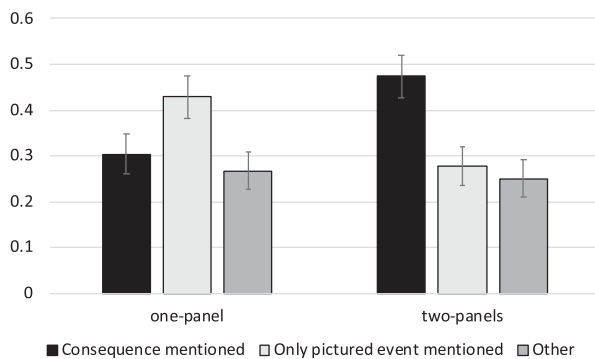


Figure 3. Event type of continuations: proportion of continuation types in the one-panel and two-panel conditions (Error bars show ± 1 SE).

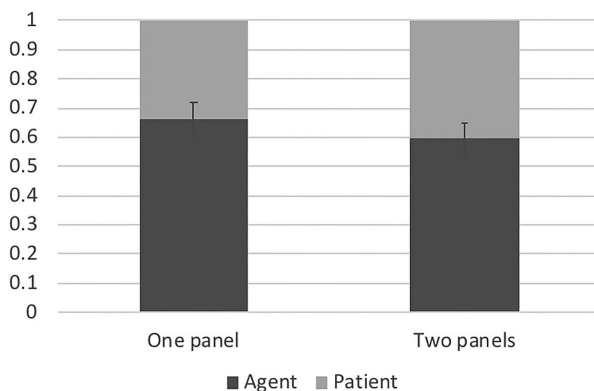


Figure 4. First-mentioned referent: proportion of continuations in the one-panel and two-panel conditions where the agent is mentioned first vs. where the patient is mentioned first (Error bars show ± 1 SE).

patient of the depicted event is mentioned first. The one- vs. two-panel manipulation has no significant effect on the choice of first-mentioned referent (glmer on agent-first continuations: $\beta = 0.073$, $SE = 0.559$, $z = 0.131$, $p = .896$).

Discussion

This experiment investigated the relation between cognitive prominence and linguistic encoding of events and entities in causal sequences. Prior linguistic work leaves open the question of whether the prominence of entities and events in mental models of discourse can be dissociated. By presenting participants with one-panel vs. two-panel segmentations of the same event, we tested whether segmentation differences boost the prominence of the patient, the consequence event, or both, as reflected in people’s choices about what is mentioned next.

The results show significant effects of panel segmentation on event prominence: Participants are more likely to mention consequences in the two-panel than in the one-panel condition. This supports the hypothesis that panel segmentation influences the prominence of consequence events. However, there is no clear support for the hypothesis that panel segmentation influences the prominence of entities, at least as reflected in continuations: The panel segmentation manipulation does not influence choice of the first-mentioned entity (agent vs. patient). (It is possible to describe a consequence without starting with the patient, e.g. ... *and then the girl apologized profusely for accidentally tripping the lady.*)

These results suggest that language users separately track expectations about who is likely to be mentioned next and what kind of event is likely to be mentioned next. If these patterns reflect cognitive prominence, this result indicates that our representation of the event prominence is separable from entity prominence. Thus, these results argue in favour of models of event representation that encode prominence in a way that allow for a dissociation between entities and the events they participate in. Broadly speaking, this study shows that visuospatial cues to event segmentation influence expectations about upcoming discourse.

Notes

1. The “other” category includes subsequent events not directly causally connected to the depicted event (e.g. “and then we went to go eat dinner”), events occurring before the depicted event (e.g. “and then Jen decided to play a trick on Jane”), etc. Continuations analyzed as

“unclear” are also in this category. Recall that the task was intentionally open-ended, to avoid biasing participants’ continuations. An unsurprising consequence is that both conditions yield approximately 25% “other” continuations. The proportion of “other” continuations in the two conditions does not differ (see Results) and thus does not impact the main claims of this paper.

2. Use of pronouns vs. NPs is not relevant: Prior work indicates this is driven by antecedent grammatical role (with pronouns preferring subjects, e.g. Fukumura & Van Gompel, 2010; Kehler & Rohde, 2013) – which we do not manipulate. Indeed, as expected, more pronouns were produced in both conditions when referring to the agent rather than the patient.
3. We excluded as uncodable those continuations where the first-mentioned referent was something other than a singular third-person expression referring to the agent/patient, because the purpose of the entity-based analysis is to assess the relative prominence of the agent and the patient.

Acknowledgements

Many thanks to David Cheng-Huan Li for drawing the stimuli. Thanks are also due to Emily Fedele, David Li, Josephine Lim, Tyler Anne Isaman, Michal Meyers and Iris Ouyang for help with data collection and analysis. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health.

Disclosure statement

No potential conflict of interest was reported by the author.

Funding

Research reported in this publication was supported by the Eunice Kennedy Shriver National Institute of Child Health and Human Development of the National Institutes of Health [R01HD061457].

References

- Ariel, M. (1990). *Accessing noun-phrase antecedents*. London: Routledge.
- Arnold, J. E. (2001). The effect of thematic roles on pronoun use and frequency of reference continuation. *Discourse Processes*, 31, 137–162.
- Chafe, W. L. (1976). Givenness, contrastiveness, definiteness, subjects, topics and point of view. In C. N. Li (Ed.), *Subject and topic* (pp. 27–55). New York, NY: Academic Press.
- Cohn, N. (2013). Visual narrative structure. *Cognitive Science*, 37(3), 413–452.
- Cohn, N. (2015). Narrative conjunction’s junction function: The interface of narrative grammar and semantics in sequential images. *Journal of Pragmatics*, 88, 105–132.
- Cohn, N. (2018). In defense of a ‘grammar’ in the visual language of comics. *Journal of Pragmatics*, 127, 1–19.
- Cohn, N., & Kutas, M. (2017). What’s your neural function, visual narrative conjunction? Grammar, meaning, and fluency in sequential image processing. *Cognitive Research: Principles and Implications*, 2(27), 1–13.
- Cohn, N., & Paczynski, M. (2013). Prediction, events, and the advantage of agents: The processing of semantic roles in visual narrative. *Cognitive Psychology*, 67(3), 73–97.
- Cohn, N., & Wittenberg, E. (2015). Action starring narratives and events: Structure and inference in visual narrative comprehension. *Journal of Cognitive Psychology*, 27, 812–828.
- Comrie, B. (1985). Causative verb formation and other verb-forming morphology. In T. Shopen (Ed.), *Grammatical categories and the lexicon. Language typology and syntactic description* (Vol. 3, pp. 309–348). Cambridge: Cambridge University Press.
- Dowty, D. (1991). Thematic proto-roles and argument selection. *Language*, 67, 547–619.
- Eisner, W. (1985/2008). *Comics and sequential art*. New York, NY: Norton.
- Fillmore, C. J. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory* (pp. 1–88). New York, NY: Holt, Rinehart and Winston.
- Fukumura, K., & Van Gompel, R. (2010). Choosing anaphoric expressions: Do people take into account likelihood of reference? *Journal of Memory and Language*, 62(1), 52–66.
- Garvey, C., & Caramazza, A. (1974). Implicit causality in verbs. *Linguistic Inquiry*, 5(3), 459–464.
- Givón, T. (1983). *Topic continuity in discourse: A quantitative cross-language study*. Amsterdam: John Benjamins.
- Grüter, T., Rohde, H., & Schafer, A. (2017). Coreference and discourse coherence in L2: The roles of grammatical aspect and referential form. *Linguistic Approaches to Bilingualism*, 7(2), 199–229.
- Gundel, J. K., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, 69(2), 274–307.
- Hartshorne, J. K. (2014). What is implicit causality? *Language and Cognitive Processes*, 29(7), 804–824.
- Hobbs, J. R. (1979). Coherence and coreference. *Cognitive Science*, 3, 67–90.
- Jasinskaja, K., Chiriacescu, S., Donazzan, M., von Heusinger, K., & Hinterwimmer, S. (2015). Prominence in discourse. In A. De Dominicis (Ed.), *Prominences in linguistics. Proceedings of the pS-prominenceS international conference* (pp. 134–153). Viterbo: University of Tuscia/DISUCOM Press.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge: Harvard University Press.
- Keenan, J. M., Baillet, S. D., & Brown, P. (1984). The effects of causal cohesion on comprehension and memory. *Journal of Memory and Language*, 23(2), 115.
- Kehler, A., Kertz, L., Rohde, H., & Elman, J. L. (2008). Coherence and coreference revisited. *Journal of Semantics*, 25, 1–44.
- Kehler, A., & Rohde, H. (2013). A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics*, 39, 1–37.
- Lascarides, A., & Asher, N. (1993). Temporal interpretation, discourse relations and commonsense entailment. *Linguistics and Philosophy*, 16, 437–493.
- Magliano, J. P., Kopp, K., Higgs, K., & Rapp, D. N. (2017). Filling in the gaps: Memory implications for inferring missing content in graphic narratives. *Discourse Processes*, 54(8), 569–582.

- Mandler, J., & Johnson, N. (1977). Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 9, 111–151.
- Mann, W. C., & Thompson, S. A. (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8(3), 243–281.
- McCloud, S. (1993). *Understanding comics*. New York, NY: Harper Perennial.
- Myers, J. L., Shinjo, M., & Duffy, S. A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language*, 26(4), 453–465.
- Olmstead, M. C., & Kuhlmeier, V. A. (2015). *Comparative cognition*. Cambridge: Cambridge University Press.
- Pearl, J. (2009). *Causality. Models, reasoning and inference*. New York, NY: Cambridge University Press.
- Postema, B. (2010). *Mind the gap: Absence as signifying function in comics* (PhD Thesis). Michigan State University.
- Radvansky, G. (2012). Across the event horizon. *Current Directions in Psychological Science*, 21(4), 269–272.
- Radvansky, G. A., & Copeland, D. E. (2000). Functionality and spatial relations in memory and language. *Memory and Cognition*, 28, 987–992.
- Radvansky, G., Krawietz, S., & Tamplin, A. (2011). Walking through doorways causes forgetting: Further explorations. *The Quarterly Journal of Experimental Psychology*, 64(8), 1632–1645.
- Radvansky, G. A., & Zacks, J. M. (2017). Event boundaries in memory and cognition. *Current Opinion in Behavioral Sciences*, 17, 133–140.
- R Core Team. (2018). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing. Retrieved from <http://www.R-project.org/>
- Rohde, H., & Kehler, A. (2014). Grammatical and information-structural influences on pronoun production. *Language, Cognition, and Neuroscience*, 29, 912–927.
- Rumelhart, D. E. (1975). Notes on a schema for stories. In D. G. Brown & A. Collins (Eds.), *Representation and understanding: Studies in cognitive science* (pp. 211–236). New York, NY: Academic Press.
- Saraceni, M. (2016). Relatedness: Aspects of textual connectivity in comics. In N. Cohn (Ed.), *The visual narrative reader* (pp. 115–128). London: Bloomsbury.
- Stainbrook, E. J. (2016). A little cohesion between friends; or, we're just exploring our textuality: Reconciling cohesion in written language and visual language. In N. Cohn (Ed.), *The visual narrative reader* (pp. 129–154). London: Bloomsbury.
- Stevenson, R., Crawley, R., & Kleinman, D. (1994). Thematic roles, focusing and the representation of events. *Language and Cognitive Processes*, 9, 519–548.
- Testelec, Y. (1998). On two parameters of transitivity. In V. Kulikov & H. Vater (Eds.), *Typology of verbal categories* (pp. 29–45). Tübingen: Niemeyer.
- van Dijk, T. A., & Kintsch, W. (1983). *Strategies in discourse comprehension*. New York, NY: Academic Press.
- von Heusinger, K., & Chiriacescu, S. (2013). The discourse structuring potential of differential object marking. The case of indefinite and definite direct objects in Romanian. *Revue Roumaine de Linguistique*, LVIII, 439–456.
- Wolff, P. (2003). Direct causation in the linguistic coding and individuation of causal events. *Cognition*, 88, 1–48.
- Zacks, J., Speer, N., Swallow, K., Braver, T., & Reynolds, J. (2007). Event perception: A mind-brain perspective. *Psychological Bulletin*, 133(2), 273–293.
- Zacks, J., & Swallow, K. (2007). Event segmentation. *Current Directions in Psychological Science*, 16, 80–84.