

Cross-Document Narrative Alignment of Environmental News: A Position Paper on the Challenge of Using Event Chains to Proxy Narrative Features

Ben Miller

Emory University

b.j.miller@emory.edu

Abstract

Cross-document event chain co-referencing in corpora of news articles would achieve increased precision and generalizability from a method that consistently recognizes narrative, discursive, and phenomenological features such as tense, mood, tone, canonicity and breach, person, hermeneutic composability, speed, and time. Current models that capture primarily linguistic data such as entities, times, and relations or causal relationships may only incidentally capture narrative framing features of events. That limits efforts at narrative and event chain segmentation, among other predicate tasks for narrative search and narrative-based reasoning. It further limits research on audience engagement with journalism about complex subjects. This position paper explores the above proposition with respect to narrative theory and ongoing research on segmenting event chains into narrative units. Our own work in progress approaches this task using event segmentation, word embeddings, and variable length pattern matching in a corpus of 2,000 articles describing environmental events. Our position is that narrative features may or may not be implicitly captured by current methods explicitly focused on events as linguistic phenomena, that they are not explicitly captured, and that further research is required.

1 Introduction

A story is not so much a unique container to be read in isolation, but a locus of potential and actual connections to other stories and their representations. Those connections enable, among other things, critical inference and empathic reading, or the sharing of a feeling from a story based on successfully imagining a character's perspective. News literacy (Fleming 2014, Hornik 2017), as taught at the Stony Brook Center for News Literacy, the Berkman Center, and the Rockefeller, Revson, McCormick, and Knight Foundations backed News Literacy Project center on teaching how to think critically about news, a skill similarly predicated on one's ability to evaluate and connect narratives. Understanding the narrative framing used to convey specific topics in journalism, such as news about the environment, may help explain aspects of audience engagement and comprehension.

Technologies like RDF and XML allow for the hardcoding of some of those connections in to a document at its moment of composition, but most stories, even those published by contemporary news organizations, lack the rigorous indexing they make possible. RDF, rather than directing semantic annotation at the moment of composition, is a model that can guide annotation projects taking place after publication (Vossen and Cybulska 2017). And while advancements in unifying annotations like Richer Event Description (O'Gorman et al. 2016) and schema based approaches (Simonson and Davis 2016) can increase the aggregation of existing automated annotation strategies or advance work to elucidate narrative schemas via document categorization, neither model explicitly addresses what narratologists like Genette (1980), Bal (1985), Ryan (1991), Bruner (1991), or Mani (2014) have identified as core attributes of narrative. Without explicitly capturing those core attributes, and without evaluating the extent to which those attributes are indirectly captured, tasks like critical reading, inference, and em-

pathic reading cannot be modeled successfully as they depend on discursive and phenomenological aspects of perspective like tempo, mood, person, tense, intentional state entailment, hermeneutic composability, and tone.

Capturing these features of narrative technique, such as the varying differential in speed between the passage of time between all the events in a chain that could have been documented (story time) and the passage of time between the events in a chain that were documented (narrative time), would, I contend, further research on the conceptual, emotional, and persuasive nature of narrative. However, that work first necessitates documenting the granular, connective tissue of narrative.

One domain in journalism that could benefit from increased sensitivity to narratological features is reporting on the environment. In the American context, a significant impediment to pro-environmental policy seem to be the opinions of the general public. Those opinions are shaped by social forces, including reportage on the environment. Identifying and demonstrating the narratological features used to frame different environmental stories such as weather and climatic events may help reveal connections among different populations and their stance towards environmental policy. A narrative, in the context of news about the environment, could be defined variously depending upon the particular question motivating the study. In regards to this study, it is comprised of events organized and related in such a way so as to allow for the recognition of features described by the above list of narratologists.

2 Blending Linguistic and Conceptual Descriptions of Narrative

The granular form these cross-document connections take, as identified by computational linguistics, are linguistic phenomena such as exophora, entity mentions when an entity is known in advance to be notable, absolute references such as to locations or measurements, and bibliographic citations. Exophora are the anaphoric, pronominal references that stretch across texts. This already high bar for meaningful cross-document annotations is even more challenging than it appears because, as with conceptual metaphors, these linguistic tokens are indicators of a phenomena, rather than the phenomena itself. The current shorthand is to identify a token that expresses an event (Ponti and Korhonen 2017), then annotate that token or set of tokens (Simonson 2016). For what we have described as concrete events in relation to human rights violations (e.g. events that have a particular physical outcome that is relatively independent of interpretation, such as an unlawful killing), the conflation of the linguistic and the conceptual is not highly problematic, as even simple string-based information retrieval of that type of information accounted for an F_1 score of 0.80 (Miller et al. 2014). Moving beyond this kind of superficial engagement to more conceptual categories of event knowledge is a task not dissimilar to engaging with conceptual metaphors.

Problematically for computational linguistics, even in projects reliant on manual annotation, inter-annotator agreement for rich metaphor annotation is relatively low when compared to annotation of other linguistic features such as entity category. Gordon (2015) found pairwise Cohen κ scores of 0.65 and 0.42 for annotations of 14 categories of conceptual metaphors in 1,450 sentences across their three annotators. In regard to narrative annotation, Finlayson (2015) returned a strict F_1 -score of 0.22 and a generous F_1 -score of 0.71 for interannotator agreement on manual labeling of Proppian functions. Implementing annotation for the full taxonomy for Genette's narrative schema (1983), an as yet unattempted task, would require annotations for four levels of distance, four narrator functions, four type of time, five levels, four speeds, three orders, and three frequencies. The results of Gordon et al.'s work, and the complexity of narratological frameworks indicates that there is a great deal of distance between the linguistic token and the conceptual framing described by narrative theory.

The misleading slippage between the notion of a token as an indicator and a token as the phenomena to be indicated is common in the literature on language processing of conceptual phenomena. For example, in the guidelines defining an event for the Automatic Content Extraction (ACE) program, the Linguistic Data Consortium states the following: "An Event is a specific occurrence involving participants. An Event is something that happens. An Event can frequently be described as a change of state. We will not be tagging all Events, but only examples of a particular set of types and subtypes" (LDC 2005). What is valuable at this point is not the particular definition of an event put forward by ACE, one oddly compliant with the Russian Formalist definition of an event from 1925 (Shklovsky 1991), but the conflation of event with the particular text that will be tagged. The TERQAS workshop productively complicated the definition of events as stative or dynamic situations occurring within a particular focal

extent (Pustejovsky et al. 2003). This definition blends the linguistic understanding necessary to develop a computational approach with a narratological approach sensitive to focalization and storytelling. Contemporary representations of these definitions, such as that developed as the Event and Situation Ontology (Segers et al. 2015), focus on linking dynamic and stative events in conceptual frames so as to enable better inferencing. The last complication we will address and one core to our work, as described in Sprungnoli and Tonelli (2016), is that each domain, e.g. history, journalism, literary studies, most likely has many competing, useful, underspecified definitions of what constitutes an event, and that those definitions have functions not addressed by an event definition driven by what is computationally tractable. Those domains address critical questions of influence, veridicality, and resonance that may require rethinking how to bridge the ontological and the linguistic.

3 Narrative structure as critical reading

Given the power of event chains as conceptual containers to further critical reading strategies like inferencing, our work aims at the cross-document coreference of meaningful chains while preserving their ontological and conceptual, rather than their linguistic instantiations.

Existing methods in similar work by Simonson and Davis (2015), Chambers and Jurafsky (2009), Bean and Rilof (2004), and Bagga and Baldwin (1999), focused in an incremental way on how to build approximations of narrative structures that included actors in various roles. For an example of how investigations into narrative schemas can facilitate critical thinking, Simonson and Davis (2015) raises the question as to why event chains extracted from articles indexed by the New York Times which were otherwise classified as “Murders and Attempted Murders,” were not, when police officers were the ones performing the shootings and killings. Aspects of each story’s focalization in relation to the schemas that incorporate those aspects, might facilitate better understanding as to why that disconnect is true. Bruner (1991) supports this reading in his exploration of canonicity and the breach; a narrative must draw on two competing pre-existing cultural tropes, a script that sets the background expectations and a breach which violates that script in ways that are often expected, but still worth telling. Retelling a visit to a restaurant evokes a canonical script but, per Bruner, is not a narrative. When the tip is a winning lottery ticket, the owner a long-lost relative, or the server spills a beverage to some later consequence, that script is breached in an expected, but interesting way.

One way in which those methods fall short of the ambitious goal of automatically capturing narrative strategies is that what they explicitly capture is not narrative, but events, semantic roles, dependencies, verbs, nouns, and scripts. Adoption of schema performance measures such as the narrative Cloze test (Chamers and Jurafsky 2008), demonstrate this, as the Cloze test effectively measures canonicity, but not breach, and is therefore more suited for measuring the accuracy of scripts than of narratives. While those methods may also indirectly capture narrative features, it is unclear at this point whether any of them do, and if so, to what extent. These limits contradict the central role posited by cognitive theories of narrative from Lewis (1978) and Johnson-Laird (1983) to Bruner (1991) and Gervas (2016). Narrative, for those theorists, functions as a means of organizing and communicating information about the world from the perspective of the narrator. As such, narrative offers a vehicle with which embedded knowledge of an environment is simultaneously interpreted and captured, then disseminated in a highly complex structure. Integrating Bruner’s 10-point schema of narrative diachronicity, particularity, intentional state entailment, hermeneutic composability, canonicity and breach, referentiality, genericness, normativeness, context sensitivity and negotiability, and narrative accrual with Genette’s would offer an ontology blending phenomenological approach and discursive approaches. In essence, it would allow for aligning event chains across documents in regard to their linguistic, phenomenological, and discursive perspectives, allowing for narrative alignment to recognize correspondences across both event as documentation of a phenomena, and event as documentation of a complex perspective.

Given the challenge of recognizing meaningful—e.g., multivariate—connections between documents, our work begins with a series of simplifying assumptions. First is that events and narrative features have meanings that are stable after they are extracted from a particular story and stored in a new format. Second, an event can be indicated by a continuous or discontinuous linguistic token or set of tokens such that while annotation is required, paraphrasing and glossing are not. Third, that a particular linguistic token that stands in for an event is commensurate with semantically similar tokens in other contexts and documents. And lastly, that the general applicability of a priori morphological frameworks like Propp’s

(1968) functions is limited despite expansions and emendations (Gilet, 1998; Gervas 2013; Finlayson 2016). Applications of Propp's framework beyond his primary corpus of Russian fairy tales from the Afanasyev collection have required extensive adjustments. Dundes (1963), in his analysis of Native American folk tales, used a similar morphological approach focused on motifs that blended the work of Propp, Kenneth Pike, and Stith Thompson. To connect with Propp's framework, he equated the disequilibrium-to-equilibrium transition common to his corpus to be equivalent to either a state of either surplus or lack. Dundes did so, as the second state, lack, motivated the stories Propp analysed. Those using Propp's functions note that his work was intended to support a multicultural analysis of folklore, but that the schema is limited in its descriptive power. Colby (1973) in his discussion of North Alaskan Eskimo folktales has to alter Propp's morphological framework to suit a different cultural tradition. Though functions are still defined by their position in a narrative sequence, he introduces two other mechanisms that define a function: sequence structures such as loops, and categorical rules operative when figures appear in certain sets. These new components of narrative he named Eidons, indicating that only five of Propp's 31 functions were relevant for his corpus. Studies like these demonstrate the extensive theoretical work necessary to generalize morphological approaches for narrative analysis. Not only were new functions necessary, but the story grammar itself, the basis of a morphological approach alongside the notion of a narrative's divisibility, required new rules.

These descriptive boundaries for schemas like Propp's show that this type of structural approach is an example of a phenomenological-discursive framework; Propp, Dundes, and Colby each described narrative with a technical structure from a limited perspective and that description is most valuable for understanding the world represented by their respective corpora.

4 Examples of narrative structure in stories about the weather

This study's preliminary empirical work began by finding concrete events using token-focused methods from four corpora drawn from contemporary news articles about the environment. The four corpora themselves each comprised 532-584 articles from 2000-2010 drawn from English-language publications with each article ranging in length from 750-1250 words. The total corpus is 2,258 articles spanning 1.96m tokens. Respectively, the corpora on earthquakes, hurricanes, pollution, and tsunami referenced terms associated with deaths or injuries 875, 334, 208, and 681 times, respectively. The references described events as varied as, "killed more than 230,000 people," and "killed 35,322 people in Sri Lanka," in an article about an earthquake and one about a tsunami, to "left at least 200 people dead or missing" in a report about a hurricane, to "oxygen starved dead zones," in an article about pollution. While inferences can be drawn about the way each of those events are covered and the human cost of natural and man-made disasters, assessing the deployment of more ontological categories of narrative understanding such as canonicity and breach to each of these fatal events is what would enable more complex, critical engagement.

As an example, consider the phrase, "dead or missing." This nomenclature only appears 21 times across the 2,258 articles about natural disasters. Of those 21 times, 14 are in articles about tsunamis, 1 in an article about a hurricane, and 6 in articles about earthquakes (of which one was about the combined earthquake-tsunami event that devastated Fukushima, Japan). In one article on the 2004 Indian Ocean tsunami, the narrative unit centered on the stative event, "dead or missing," proceeds as follows. A local official is quoted as recounting that victims' bodies have been identified, but his knowledge is put forward as lacking in relation to the national origin of those individuals. One country is put forward as particularly hard-hit with an explanation provided for why, and an official from the foreign country is quoted enumerating the toll, indicating the method of identification, and suggesting that the work of identification and repatriation will be more difficult from here on. Comparing the elements of this article's narrative versus the narrative elements of other articles describing similar events has the potential to reveal, for example, both aspects of canonicity and breach, and issues of what Bruner calls hermeneutic composability. That term refers to the range of possible interpretations supported by a given set of phrases, a range that journalism, and computational approaches suggest we limit by selecting the most common interpretation. Aligning narratives about similar events within a corpus would allow for the evaluation of that process, one Bruner refers to as narrative banalization. While not a feature of a narrative per se, evaluating the relationships between narratives and the events they purport to convey is a core task for the computational analyses of narrative.

5 Conclusions and future work

Based on work like Simonson (2018), Caselli and Vossen (2017), and Spurgnoli and Tonelli (2016) that focus on linguistic event chain extraction, and work like that reported by Minard, et al (2015) and Miller (2015, 2015b) on event cross-document coreference, we presume that chains of discursive and descriptive events can be captured, albeit with difficulty at present (Laparra et al 2017). These chains would explicitly include the dynamic events like quoting, recounting, and describing, along with the dynamic events like identification, repatriation, and enumeration, along with stative events like dead and missing. A current challenge is that although these methods would appear to capture aspects relevant to Bruner's phenomenological framework, they do so idiosyncratically. For example, Simonson's Chambers and Jurafsky-inspired schemas can find frequently occurring narrative patterns in news corpora, but cannot address aspects like the canonicity of narrative structures and their breach. In stories of the 2004 Indian Ocean earthquake and tsunami, stories frequently exemplifying the canonical form of a natural disaster and the recovery of first world citizens in second or third-world contexts, the form is breached when the hard reality of the only possible method of identification, dental records, eclipses the particular impact of the event on a given first-world nation far from the event's epicenter.

At the heart of this research is an effort to understand the relationship among narrative form and audience engagement with complex stories in the context of journalism. It is with narratological details such as those described above that we might find answers to questions such as how do different weather events get covered, why do audiences respond to some coverage but not others, what kinds of canonical frames are relied upon for covering different types of weather events and disasters, and how the coverage of events might better align with the significance of those events. While current methods of schema extraction and narrative alignment are remarkable in their ability to identify canonical scripts, they may only accidentally capture aspects of narrative necessary for the answering of complex questions. Future work needs to go beyond measures of accuracy to canon such as the Narrative Cloze test to instead assess the relationship of canonicity and breach, and to quantify the extent to which narrative features, even ones as culturally and generically limited as the Proppian functions, are represented in annotations or extractions.

Our future work will be to conduct empirical research on our weather-event corpora to assess the efficacy of, first, cross-document narrative alignment using fuzzy matching of word embedding-based event chains, and second, to manually assess the extent to which that method indirectly captures narrative features concordant with a blended Bruner-Genette framework. This position paper, by describing how that framework would support critical reading, interpretation, and inference of events in the news, makes the argument that events must be considered more as conceptual containers that incorporate the commonplace with its exception, than as linguistic tokens reflective of the common.

Acknowledgements

The author would like to thank Ayush Shrestha, Jennifer Olive, Yanjun Zhao, and Shakthidar Gopavaram, collaborators on prior research into narrative segmentation and alignment upon which this work builds. I would also like to thank the reviewers of this paper for their valuable feedback, in particular about morphological approaches to the narrative analysis of folktales, and the ongoing challenge of capturing narrative sequences.

References

- Bagga, A., and Baldwin, B. (1999). "Cross-document event coreference: Annotations, experiments, and observations." *Proceedings of the Workshop on Coreference and its Applications*. Association for Computational Linguistics.
- Bal, M. (1985). *Narratology: Introduction to the Study of Narrative*. Trans. Christine van Boheemen. Toronto: U of Toronto P.
- Bean, D. L., & Riloff, E. (2004). Unsupervised Learning of Contextual Role Knowledge for Coreference Resolution. In *HLT-NAACL* (pp. 297-304).
- Bruner, J. (1991). The narrative construction of reality. *Critical inquiry*, 18(1), 1-21.

- Caselli, Tommaso, and Piek Vossen. "The Event StoryLine Corpus: A New Benchmark for Causal and Temporal Relation Extraction." *Proceedings of the Events and Stories in the News Workshop* (EventStory 2017). 2017
- Chambers, N., & Jurafsky, D. (2008, June). Unsupervised Learning of Narrative Event Chains. In *ACL* (Vol. 94305, pp. 789-797).
- Chambers, N., & Jurafsky, D. (2009). Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2* (pp. 602-610). Association for Computational Linguistics.
- Dundes, A. (1963). Structural Typology in North American Indian Folktales. *Southwestern Journal of Anthropology*, 19(1), 121-130.
- Finlayson, M. A. (2015). Propplearner: Deeply annotating a corpus of Russian folktales to enable the machine learning of a Russian formalist theory. *Digital Scholarship in the Humanities*, fqv067.
- Finlayson, M. A. (2016). Inferring Propp's functions from semantically annotated text. *Journal of American Folklore*, 129(511), 55-77.
- Fleming, J. (2014). Media literacy, news literacy, or news appreciation? A case study of the news literacy program at Stony Brook University. *Journalism & Mass Communication Educator*, 69(2), 146-165.
- Genette, G. (1980). *Narrative discourse* (JE Lewin, Trans.). Ithaca, NY: Cornell UP.
- Genette, G. (1983). *Narrative discourse: An essay in method*. Cornell University Press.
- Gervás, P. (2016). An Exploratory Model of Remembering, Telling and Understanding Experience in Simple Agents. In *Proceedings of the Workshop on Computational Creativity, Concept Invention, and General Intelligence (C3GI 2016)* Vol-1767. CEUR Workshop Proceedings.
- Hornik, R. (2017). Why News Literacy Matters: A New Literacy for Civil Society in the 21st Century. Digital Resource Center.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness* (No. 6). Harvard University Press.
- Laparra, E., Agerri, R., Aldabe, I., & Rigau, G. (2017). Multi-lingual and Cross-lingual timeline extraction. *Knowledge-Based Systems*, 133, 77-89.
- Gilet, P. (1998). *Vladimir Propp and the Universal Folktale: Recommissioning an Old Paradigm--story as Initiation* (Vol. 17). Peter Lang Pub. Incorporated.
- Gordon, J., Hobbs, J.R., May, J., et al. (2015). "A Corpus of Rich Metaphor Annotation". In *Proceedings of the Third Workshop on Metaphor in NLP*.
- LDC. (2005). Ace (automatic content extraction) English annotation guidelines for events ver. 5.4.3 2005.07.01. Technical report, Linguistic Data Consortium
- Lewis, D. (1978). Truth in fiction. *American Philosophical Quarterly*, 15(1), 37-46.
- Mani, I. (2014). Computational narratology. *Handbook of narratology*, 84-92.
- Miller, B., Shrestha, A. & Subtirelu, N. (2014) "NLP Approaches to Rights Violation Classifying." Paper presented at the annual *RightsCon meeting*, San Francisco, California.
- Miller, B., Shrestha, A., Olive, J., & Gopavaram, S. (2015). Cross-Document Narrative Frame Alignment. In *OASIS-OpenAccess Series in Informatics* (Vol. 45). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Miller, B., Olive, J., Gopavaram, S., et al. (2015b). A method for cross-document narrative alignment of a two-hundred-sixty-million word corpus. In *Big Data (Big Data), 2015 IEEE International Conference on* (pp. 1673-1677). IEEE.
- Minard, A. L., Speranza, M., Agirre, E., Aldabe, I., van Erp, M., Magnini, B., ... & Urizar, R. (2015). Semeval-2015 task 4: Timeline: Cross-document event ordering. In *9th International Workshop on Semantic Evaluation (SemEval 2015)* (pp. 778-786).
- O'Gorman, T., Wright-Bettner, K., & Palmer, M. (2016). Richer Event Description: Integrating event coreference with temporal, causal and bridging annotation. *Computing News Storylines*, 47.

- Gervás, P. (2013). Propp's Morphology of the Folk Tale as a Grammar for Generation. In *Proceedings of the 2013 Workshop on Computational Models of Narrative*, Dagstuhl, Germany
- Ponti, E. M., & Korhonen, A. (2017). Event-Related Features in Feedforward Neural Networks Contribute to Identifying Causal Relations in Discourse. *LSDSem 2017*, 25.
- Pustejovsky, J., Hanks, P., Sauri, R. et al. (2003, March). The timebank corpus. In *Corpus linguistics* (Vol. 2003, p. 40).
- Ryan, M. L. (1991). *Possible worlds, artificial intelligence, and narrative theory*. Indiana University Press.
- Segers, R., Vossen, P., Rospocher, M., Serafini, L., Laparra, E., & Rigau, G. (2015). Eso: A frame based ontology for events and implied situations. *Proceedings of Maplex 2015*.
- Shklovsky, V. (1991) *Theory of Prose*, trans. B. Sher, Intr. G.L. Burns. Elmwood Park, IL: Dalkey Archive Press.
- Simonson, D., & Davis, A. (2016). NASTEA: Investigating Narrative Schemas through Annotated Entities. *Computing News Storylines*, 57.
- Simonson, D., Davis, A. R., & Solutions, E. (2015). Interactions between Narrative Schemas and Document Categories. *ACL-IJCNLP 2015*, 1.
- Sprugnoli, R., & Tonelli, S. (2016). One, no one and one hundred thousand events: Defining and processing events in an inter-disciplinary perspective. *Natural Language Engineering*, 1-22.
- Propp, V. (1968). *Morphology of The Folk Tale* (second edition). University of Texas Press.
- Vossen, P., & Cybulska, A. (2017). Identity and Granularity of Events in Text. *arXiv preprint arXiv:1704.04259*.