

Narrative Theory for Computational Narrative Understanding

Andrew Piper

Department of Languages, Literatures, and Cultures
McGill University
andrew.piper@mcgill.ca

Richard Jean So

Department of English
McGill University
richard.so@mcgill.ca

David Bamman

School of Information
UC Berkeley
dbamman@berkeley.edu

Abstract

Over the past decade, the field of natural language processing has developed a wide array of computational methods for reasoning about narrative, including summarization, common-sense inference, and event detection. While this work has brought an important empirical lens for examining narrative, it is by and large divorced from the large body of theoretical work on narrative within the humanities, social and cognitive sciences. In this position paper, we introduce the dominant theoretical frameworks to the NLP community, situate current research in NLP within distinct narratological traditions, and argue that linking computational work in NLP to theory opens up a range of new empirical questions that would both help advance our understanding of narrative and open up new practical applications.

1 Introduction

Research in NLP has seen an increasing attention to narrative understanding over the past decade. Indeed, the NLP community is not alone. From studies in economics (Shiller, 2020) to climate science (Bushell et al., 2017) to political polarization (Kubin et al., 2021) to mental health (Adler et al., 2016), there is a growing consensus that narrative is a key concept for understanding human behavior and beliefs. Narrative is a core mechanism through which human beings come to understand their world, find meaning, motivate their actions and those of others, and create communities. As narratologists often highlight, narrative is a universal practice among all human cultures across all time periods (Barthes and Duisit, 1975).

In developing computational methods to understand narrative and its various social, personal, and

cultural functions, the NLP community has drawn on a wide range of theoretical perspectives on narrative (both implicit and explicit), without the benefit of situating those perspectives within a broader organizing theoretical structure. Such an organizing structure can encourage terminological consistency and methodological clarity in terms of research goals, while also illuminating the points of contact between seemingly unrelated research agendas.

This position paper seeks to provide the beginnings of such a unifying framework for the computational study of narrative. Drawing on the multiple expertises of literary and textual studies on the one hand and natural language processing on the other, our aim is to provide an overview of the different theoretical components of narrative study and highlight promising NLP research in those domains. By providing the NLP community with a richer and more coherent theoretical foundation for implementing computational solutions for narrative understanding, our goal is to bring together otherwise disconnected lines of inquiry as well as reduce competing conceptual frameworks that can inhibit progress. In doing so, the field will be better positioned to address the major research challenges that we describe in the closing section.

In what follows we organize our paper in the following way: first, we provide a general definition of narrativity, integrating existing theoretical frameworks; second, we review the fundamental components of narrative understanding and the research opportunities within each domain; third and finally, we identify a few salient higher-level problems of narrative understanding for the NLP community to address as they relate to questions of social change, creative industries, and social and mental well-being.

2 The elements of narrativity

2.1 Towards a working definition of narrativity

Interestingly, there is no consensus when it comes to a minimal definition of whether a sequence of tokens can be classified as narrative. Many recent theories, however, do tend to converge around notions of time and process (Ricoeur, 2014; Herman, 2009; Walsh, 2018; Sternberg, 1992). As Herman (2009) writes, “Narrative roots itself in the lived, felt experience of human or human-like agents interacting in an ongoing way with their surrounding environment . . . Narrative, in other words, is a basic human strategy for coming to terms with time, process, and change” (21).

Such definitions are consistent with a large body of Aristotelian-inspired narrative theory that emphasizes “change of state” as an essential component of narrative (Liveley, 2019). Narrative is defined by sequences that represent a transformational experience or an expression of “disequilibrium” (Bruner, 1991). These include Freytag (1895)’s pyramid, Van Dijk (1976)’s “Problem-Attempt-Outcome” model, and Sternberg (1992)’s model of “Suspense-Curiosity-Surprise.”

Common to all of these definitions is an attention to “story-level” phenomena (also called *diegesis*), i.e., the structure of events that unfold in a narrative. Genette (1983) is the work best known for introducing an emphasis on the perspectival nature of narrative, i.e., that narratives must have narrators. Narratives not only encode time, but also point of view in the ordering of information. In what has come to be known as the classical model (Fig. 1), Genette introduces three primary relational terms: *story* refers to the events recounted, *discourse* to the order and economy of their telling, and *narrating* to the narrator’s role in shaping this information.¹

Genette then introduced three further terms to capture the relationship between these dimensions, *tense*, *mood*, and *voice*. Genette extrapolates from the linguistic meanings of these terms to capture specific narratological features. These include aspects of time and the ordering of events (*tense*); the relationship between eventfulness and description (*mood*); and aspects related to perspective, such as point of view, dialogue, and focalization (*voice*).

The post-classical intervention in narratology, often dated to the 1990s, emphasized that “narra-




Figure 1: Genette (1983)’s narrative triangle.

tivization is at once a textual property and a cognitive process” (Liveley, 2019). The key concept introduced here is that of “situatedness,” i.e., narrative representation is “situated in—must be interpreted in light of—a specific discourse context or occasion for telling” (Herman, 2009). *Situatedness* foregrounds the role that audiences and context play in shaping narrativity and can refer to a variety of social contexts, including medium (oral, visual, textual), platform (Facebook, Twitter, Reddit), community type, and cognitive processes (Herman, 2009). Understanding the situatedness of narratives also opens the door to understanding the interactional nature of storytelling introduced by Georgakopoulou (2007)’s theory of “small stories,” where narratives are not standalone objects but can occur through time, such as Facebook status updates (Page, 2010).

Thus we could model narrative theory as engaging with two levels of interactions (Fig. 2), where classical theory focuses on interactions between different narrative features and postclassical theory on audience interactions with feature level interactions. It is this interactional understanding of narrative that has led post-classical narratologists to speak of “degrees of narrativity” (Giora and Shen, 1994; Herman, 2009; Pianzola, 2018). Rather than think of narrative as a binary class, “narrativity” represents a scalar construct that captures the dynamic interactions among narrative features and narrative situations (Pianzola, 2018).

Given this framework, it remains necessary to explicate the elementary features that these higher level interactions might consist of and how we can observe them. Accordingly, we propose the following minimal model of narrativity, which highlights the elements that must be present in order for a symbolic sequence to aspire to narrativity. While we do not expect all variables will be explicit in any

¹*Story* and *discourse* are English-language translations of the earlier Russian formalist terms *fabula* and *syuzhet*.




Figure 2: Levels of narrative analysis.

given sequence, they are all implicitly necessary for narrativity to occur. Following the “degrees of narrativity” hypothesis, we expect that recipients’ response to narrativity would increase as the explicitation of these variables increases.

According to our minimal definition, narrativity can occur when:

A	Someone
B	tells
C	someone
D	somewhere that
E	someone
F	did something(s)
G	[to someone]
H	somewhere
I	at some time
J	for some reason

Thus what defines the presence of narrativity are the eight features: A. teller, B. mode of telling, C. recipient, D. situation, E. agent, F. one or more sequential actions, G. potential object, H. spatial location, I. temporal specification, J. rationale.

In what follows, we organize our sections according to these features, drawing attention to their further elaboration through existing theoretical frameworks and ongoing NLP work in each area. We close by exploring how better understanding their interactions can lead to further insights into major questions related to narrative understanding.

2.2 Agents (E, G)

Where classical narratology invests more heavily in the perspectival nature of narrative—the act of telling—post-classical models have re-ignited a focus on the centrality of characters or “agents” (Frow, 2014; Jannidis, 2009; Eder et al., 2010; Palmer, 2004), first explored by Propp (2010). For example, Fludernik (2002) argues that a concept of “experientiality” is definitive for narrativity: “Expe-

rientiality ... reflects a cognitive schema of embodiedness that relates to human existence and human concerns. ... In my model there can be narratives without plot, but there cannot be any narratives without a human (anthropomorphic) experiencer of some sort” (9). The nature and distribution of agents thus marks an essential element of computational narrative study. Within NLP, agents have been centered both in their detection and in predicting the relations between them.

Agent detection. While not always framed this way, agent detection can thus be seen as a foundational aspect of computational narratology. In order to better understand the role of agents within narratives, we need reliable systems to extract and identify them. Work in agent-focused NLP has emphasized expanding our understanding of characters beyond *named* entities, as in agents such as “the coachman” or “the frog,” through the concept of animacy detection (Vala et al., 2015; Jahan et al., 2018; Coll Ardanuy et al., 2020; Karsdorp et al., 2015), while also differentiating characters from other named referents (Jahan and Finlayson, 2019). Other work has utilized annotated data to improve NER performance in the literary domain (Bamman et al., 2019), making sizable differences in the accuracy of detecting such entities.

The ability to recognize agents with increasing accuracy opens up a number of research questions on their role within narratives that aligns with longstanding theoretical frameworks. For example, Propp (2010) theorized that character is defined as a limited set of narrative “functions,” while Forster (1985) introduced the distinction between round/flat characters with respect to their emotional depth. More recent work has focused on representations of gender (Cheng, 2020; Kraicer and Piper, 2019; Underwood et al., 2018; Piper, 2018a), religion (Terman, 2017), psychology (Rashkin et al., 2018; Brahman and Chaturvedi, 2020), desire fulfillment (Chaturvedi et al., 2016a) and power (Sap et al., 2017) with respect to narrative agents.

This work is driven both by the ability to recognize agents within text and ascribe attributes to them. While the former has seen progress in NLP, **character attribute inference** is an open research area that requires more attention: how do we best theorize attributes like gender, ethnicity, religion, emotion, function and power in agents within narrative texts, and build models for their estimation that respect the biases present in the data (and not those

external to it)? Better attribute inference can facilitate research to identify what [Blodgett et al. \(2020\)](#) call “harms of representation,” i.e., how different communities are underrepresented or stereotypically portrayed, which has a strong connection to theories of personal dignity ([Honneth, 1996](#)).

Relation detection. A second important tradition for understanding narrative agents is related to the field of social network theory. Character “schemas” have played an important role in narrative understanding, including [Propp \(2010\)](#)’s functions, protagonist/antagonist relations derived from classical tragedy ([Moretti, 2013](#)), [Girard \(2020\)](#)’s theory of mimetic desire, and [Woloch \(2009\)](#)’s distinction between major and minor characters.

This work depends on the accurate extraction not simply of a narrative’s agents, but of the relations between them. How do we know that two characters are “connected” or “interacting”? Within NLP, this structure has been inferred through a variety of methods (see [Labatut and Bost \(2019\)](#) for an overview), including the use of explicit quotations ([Elson et al., 2010](#)), though there remains a great deal of ambiguity in inferring the connections between characters (in scenarios such as thinking about, looking at, passing by, etc.). More recent work has examined the characterization of such relationships, both in supervised models that presume a fixed set of relations, such as positive/negative ([Chaturvedi et al., 2016b](#)) and familial/professional ([Makazhanov et al., 2014](#); [Massey et al., 2015](#)), and in unsupervised models that learn a set of descriptors ([Iyyer et al., 2016](#)).

While work in NLP has focused on the task of relation detection, a great deal of work remains to allow such methods to be usable in practice. One of the most challenging yet potentially transformative goals is **character relationship prediction** for family relations. Identifying which character pairs are siblings or parents/children is very difficult in practice, where such information is often expressed through a single statement or is altogether left to be inferred. The ability to estimate familial structures in narrative would enable a host of research questions related to this fundamental human social unit, such as: How have family relations changed over time, especially with respect to questions of size, hierarchy, and conflict? How has the agency of children evolved, either in imaginative depictions or within clinical and/or judicial settings?

2.3 Events (F)

As can be seen from the definition of narrativity in §2.1, agent-based actions form the core of the diegetic universe of a narrative (i.e., what happened). We can differentiate between two classes of problems this raises: understanding the ordering of events (and any discrepancies between *story* and *discourse*) and classifying events into more general functional units within a narrative.

Narrative sequence. Narrative *discourse* is comprised first and foremost of a sequence of events that are selected and ordered by a narrator. This order may correspond more or less closely and with more or less selectivity with the underlying events of the *story*. That is to say, given a more or less infinite set of possible events to report about the world, a narrator will choose a selection of events to report and order those events more or less sequentially with their occurrence in the storyworld.

While much work in NLP has focused on the core problem of event detection ([Nguyen and Grishman, 2015](#); [Chen et al., 2015](#); [Nguyen et al., 2016](#); [Feng et al., 2016](#); [Ahn, 2006](#); [Li et al., 2013](#); [Yang and Mitchell, 2016](#))—including event detection within literature ([Sims et al., 2019](#))—a range of work has also considered assembling sequences of events into stereotypical frames, modeling the linear order of events (along with their participants) into narrative event chains ([Chambers and Jurafsky, 2008, 2009](#)) probabilistic frames ([Cheung et al., 2013](#)) and schemas ([Chambers, 2013](#)).

In applying the construct of “scripts” ([Schank and Abelson, 1977](#))—the experiential relatedness of event sequences in discourse—this work allows, first, for an improved understanding of **eventfulness** ([Hühn, 2014](#)), i.e., the ratio of events relative to the amount of description and/or expected events (for which Plato used the terms *diegesis* (events) and *mimesis* (description)), allowing for better estimates of narrative categories like “pacing” and “suspense” ([Sternberg, 1992](#); [Doust and Piwek, 2017](#); [Wilmot and Keller, 2020](#)). Additionally, knowledge of event sequences can be used to better understand the **temporal relations** between events (which we discuss in the next section), as well as the logical relations between events with respect to questions of **causality** ([Mostafazadeh et al., 2016](#)). Causal reasoning has been identified as one of the primary functions of narrative ([Todorov, 1981](#); [Graesser et al., 2002](#)), and thus better modeling causal relations in stories can help

researchers understand the representation of human reasoning within them.

Narrative structure. The ability to resolve narrative events into higher-level abstractions represents another important area of research for narrative understanding, which we group into passage- and document-level problems.

Passage-level challenges in narrative structure include the identification of **narrative levels** (Reiter et al., 2019), where narratives include other narratives within them (cf. *1001 Nights* as a classic example) and have complex interactions with point of view (described in §2.6 below). A related problem is **narrative scene detection**, which attempts to observe the spatial, temporal, and agential boundaries between story segments. While “text worlds” (Mikhalkova et al., 2020) and “narremes” (Delmonte and Marchesini, 2017) have been proposed as alternate terms, we recommend the use of *scenes* to differentiate such horizontally structured boundary-detection issues from the vertical ones of narrative *levels* as well as the higher-level problem of **narrative plotline detection**, i.e., the act of assembling scenes and levels into more general narrative units defined by agents who may range over both time and space (Wallace, 2012). Doing so, we arrive at the following hierarchical scheme: *event-scene-level-plotline-plot*.

At the document level, work in NLP has relied on a number of different theoretical frameworks for categorizing narrative segments into higher-level types. Ouyang and McKeown (2014), Swanson et al. (2014), Levi et al. (2020) and Saldias and Roy (2020) all leverage the categorization outlined by Labov and Waletzky (1967), focusing especially on the “most reportable event (MRE)” as an indicator of a narrative’s more general meaning. While valuable for the analysis of “small stories” (Georgakopoulou, 2007), MRE is a limited construct for addressing longer, more complex narratives. Instead a preferred focus should attempt to explore the theoretical literature’s focus on “dis-equilibrium” and “change of state” (Bruner, 1991; Herman, 2009). Papalampidi et al. (2019) operationalizes the concept of “turning point” identification based on how-to guides for screenwriters (Hauge, 2017), Piper (2015) uses a similar construct for novels based on models of religious conversion, and a number of studies implement Kurt Vonnegut’s theory of the plot arc as the narrative discourse of rising and falling “fortune,” often ap-

proximated as “sentiment valence” (Jockers, 2015; Reagan et al., 2016; Boyd et al., 2020). Considerably more work needs to be done to understand the formal and cognitive conditions of such narrative pivots.

Also important to emphasize here is that no single approach is universally appropriate to understanding narrative structure. Much future work remains in exploring further ways of modeling narrative types and validating existing models more thoroughly. We take up the relevance of this issue more fully in §3.1.

2.4 Temporality (I)

Narrative theory organizes the category of time at its most general level according to the binary scheme of “narrative time” versus “narrated time,” where the latter refers to the amount of time passing within the *story* and the former refers to the amount of time passing within *discourse*. One can tell a “short-lived” story over hundreds of pages (e.g., James Joyce’s *Ulysses*) or a “long-lived” story in just a few lines (e.g., Biblical stories of human generations). Thus the first level of temporal analysis for narrative must aim to understand the relationship between the “time-of-telling” (narrative time) and the “time-of-what-is-told” (narrated time).

Underwood (2018) provides an empirical approach to this concept, annotating the temporal duration (narrated time) of passages of 250 words (i.e., within a fixed window of discourse). While computational work on **elapsed narrated time prediction** has begun (Yauney et al., 2019), much work remains to be done. This work has the potential not only to open up new empirical studies into pacing and **suspense detection** (Doust and Piwek, 2017), which theorists assume is an important dimension in audience response to narrative (Sternberg, 1992). It can also provide insights into changing historical dispositions to narrative temporality to better understand large-scale questions of cultural evolution.

A second level of temporal analysis focuses on questions of “anachrony,” the misalignment of discourse sequence from story sequence (i.e., analepsis and prolepsis or flashback and flashforward). Genette (1983) again provides a theoretical scaffolding for this work, enumerating a range of fine-grained categories of anachrony, and providing examples of scene-level annotations for such works as Homer’s *Iliad* (for which the classical trope of beginning *in medias res* necessitates flashbacks)

and Proust’s *In Search of Lost Time*.

Within NLP, this line of research is taken-up most extensively in research related to the TempEval task (Verhagen et al., 2009; UzZaman et al., 2013), which orders events in text relative to each other, but typically operates on documents with relatively short time scales (not novels spanning decades). Much work on **temporal ordering for long documents** is required, including alternative problem definitions (such as work by Kim et al. (2020) on predicting the absolute time of day for each sentence in novels, or Kearns (2020) on labeling time disruptions).

Accurately estimating the relative temporal ordering for events, sentences and scenes in narratives has the potential to shed insight on a range of narrative questions: temporal ordering is a precondition for representations of plot; it is a requirement for any kind of causal understanding (Mostafazadeh et al., 2016); and even holistic markers such as the rate of analepsis (nostalgia) can be an indicator of cultural prestige (Piper and Portelance, 2016) or potentially mental dispositions of patients or political communities.

2.5 Setting (H)

The setting of a narrative defines where the action takes place—the physical universe where the plot unfolds, which may be grounded by referents in the real world (e.g., New York City in Truman Capote’s *Breakfast at Tiffany’s*) or entirely fictional (Hogwarts Academy in Rowland’s *Harry Potter*). In the same way that we understand an agent through the actions they perform, theoretical work also points out a similar co-determinative relationship between action and setting (Burke, 1969): the setting (such as a kitchen) influences the kinds of activities that are performed there (acts of cooking, eating), and we understand that setting through those actions.

Modeling setting naturally requires the identification of place, and much work in NLP has focused on the core tasks of entity recognition and toponym resolution in the context of narrative (DeLozier et al., 2016; Brooke et al., 2016; Bamman et al., 2019), leading to models that are able to recognize both named locations, facilities and geopolitical entities (*New York*, *Hogwarts*) as well as common ones (*the kitchen*, *the village*).

The existence of such tools has already established setting as a main area of focus within cultural analytics. Wilkens (2013) has shown how U.S.

literature was much less focused on New England in its origins than theorists have claimed and continues to be much more nationally centered than, for example, UK literature, confirming European assumptions about the provincialism of U.S. literature. Erlin et al. (2021) show how minor European literatures are subject to different stylistic pressures with respect to national setting than major languages, contradicting earlier theories.

Existing work has thus already leveraged measurements of named locations to make important theoretical advances (attesting to the value of this dimension of narrative), but there remains a range of work that can only be realized by more complex measurements of the *structure* of places in text. In this view, literary theorists have emphasized the important distinction between “space” and “place” when it comes to setting, where the latter refers to concrete geo-locations or physical settings and the former to experiential constructs (Lefebvre and Nicholson-Smith, 1991; Piatti and Hurni, 2009). Understanding narrative “space” refers to relational constructs (i.e., events are dispersed or concentrated around single/many points) as well as questions of navigability and intelligibility (as in Kafka’s labyrinthine bureaucratic or governmental spaces (Piatti et al., 2013)). Similar to the logical relations between narrative sentences (temporal or causal), narratives also encode a spatial logic that can be more or less explicit.

We can ground this theoretical perspective in a number of concrete computational tasks. As noted above, inferring plotlines (and plot more generally) requires grounding actions in the places where they unfold (so that one plotline in *Lord of the Rings* is the journey by Frodo and Sam from The Shire to Mount Doom). While NER provides the capacity to identify mentions of places, one open challenge is resolving different mentions of the same place to the same physical entity. While this is relatively simple for named locations such as *Mount Doom*, coreference resolution systems struggle with **long-document coreference** of common entities like *the house* and *the room* (Bamman et al., 2020; Toshniwal et al., 2020), which define the center of the narrative universe for many fictional constructs. Likewise, many questions can only be answered by measuring the relationship *between* places mentioned in text: how far do Frodo and Sam walk on their journey? Cultural historians have long argued for an association between mobility and agency (Sen,

1993). Systems that allow better inference on questions of spatial movement within narratives could provide important insights into changing power dynamics with respect to different social communities (e.g. how far are women and children allowed to go, either in personal life stories, real-world news accounts, or fictional mythologies?).

2.6 Perspective (A, B, C)

We have so far focused on narrative elements that fall below our double line in §2.1, i.e., that concern the diegetic universe of a narrative. In this section, we look more closely at the elements that fall above our double line, which aligns with Genette’s concept of *narrating*.

At its most general level, *narrating* refers to the perspective from which a story is told, “the generating instance of narrative discourse” (Genette, 1983, 213). It thus captures an intending consciousness (whether real or imagined) as the source of language variation in a narrative. This gives rise to a number of different research directions, including point of view, focalization and dialogue.

Point of view. The simplest form of narrative perspective pertains to the relationship between the narrator and the story, i.e., whether the narrator is an agent external to or present within the diegetic universe of the narrative. We refer to this using Bal and Van Boheemen (2009)’s distinction between external narrators (EN) and character-bound narrators (CN), colloquially understood as first- and third-person narratives. Computational models of point of view have found strong differentiators in simple features like the use of personal pronouns (Eisenberg and Finlayson, 2016), and have focused on identifying extradiegetic moments in personal narratives using subjectivity as a guide (Sagae et al., 2013). Genette (1983, 248) posits a more complex set of relations between layers of narrative involvement by narrating agents that is worth further study.

Focalization. The second level called “focalization” refers to the organization of narrative language around an experiencing agent’s consciousness within the narrative that is separate from the primary narrator. This occurs when we see events explicitly refracted through the eyes, mind, and feelings of a particular agent or character. Understanding the boundaries between the narrator’s “voice” and that of a character’s perspective is a challenging interpretive task for human readers as well as machines. “Free indirect discourse” is the

term narratologists use to denote text spans where a reader cannot reliably distinguish between two distinct perspectives. While some computational efforts have attempted to **identify free indirect discourse** (Brunner et al., 2020; Long and So, 2016), it remains an extremely challenging recognition task. One way to frame the task of focalization detection is as a subset of agent attribute detection.

Dialogue. The final elementary aspect of perspective, dialogue, represents an explicit attribution of “point of view” to a particular agent within the narrative (including indirect discourse). Dialogue attribution is a well-studied, yet still challenging problem within NLP (He et al., 2013; Muzny et al., 2017b). While computational work has examined the ways in which dialogue differs from narrative (Muzny et al., 2017a) as well as functions as a main predictor of fictional versus non-fictional storytelling (Piper, 2018b), less attended to are the ways in which dialogue can provide insights into agent-specific sub-narratives. Agents are often depicted as having different voices that can be distinguished empirically (Vishnubhotla et al., 2019); in the same way that all narratives represent a particular point of view, dialogue can reveal intra-narrative perspectives and biases. Reconstructing the degree of partiality of any given agent’s perspective within the larger narrative in which they are embedded represents an important line of future research into narrative perspective—one which requires not only high-quality models for **speaker attribution**, but also models for **inferring speakers’ attitudes** on the narrated world expressed through dialogue.

3 Scaling up narrative understanding

Our aim so far is to align existing work in NLP narrative understanding with theoretical frameworks in narratology. In this final section, we want to identify three core challenge areas that tackle larger questions regarding the social impact of narrative and that depend upon the successful integration of the different strands of NLP research outlined here. Broadly, each of these areas attempts to understand narrative as an important basis for human behavior.

3.1 Narrative economies

One of the core beliefs in narratology is that narratives are deeply informed by prior narrative structures (Genette, 1997) and social context (Herman, 2009). In order to better understand how narratives influence each other or are influenced by different

social contexts, we first need more robust systems for classifying different narrative types.

Fields as diverse as folkloristics and narrative psychology provide potential models of how to move forward. Two of the best known examples of manually constructed classification systems for fictional narratives are Thompson's Motif-Index of Folk-literature (Thompson, 1989) and the Aarne-Thompson-Uther (ATU) Tale Index. In the field of narrative psychology, Adler et al. (2016) have provided a taxonomy of life narratives that correlate with life satisfaction surveys. While these models have seen only limited computational implementation (Broadwell et al., 2018), both offer examples of models for future narrative classification.

The inter-disciplinary computational implementation of such large-scale registers of story types represents a major research challenge. If successful, it would provide invaluable insights into questions of cultural evolution, narrative influence, mental health, and transnational cultural transfer. For example, how do real-world events like wars and pandemics modulate how humans invent or tell stories? When do we see new narrative forms emerging due to different historical or personal contexts? How do information economies like social media modulate the production of narrative structures circulating in the news in real time (as with events like Black Lives Matter)? And what can narrative understanding do to further combat problematic conspiracy theories or misinformation campaigns? These are just a few of the large-scale questions that researchers will be able to address as our understanding of the nature of narrative types and their interactions improves.

3.2 Narrative responses

Considerable work exists to study the efficacy of narrative discourse when it comes to persuasion (Dal Cin et al., 2004; Braddock and Dillard, 2016) and building trust (Kubin et al., 2021). Political scientists and media studies scholars in particular have focused on the concept of "framing" to analyze how the news media shapes public attention and policy debates by deciding what aspects of an issue to emphasize or exclude in reporting events (Boydston, 2013; Card et al., 2015, 2016; Field et al., 2018).

A more sophisticated computational modeling of narrative types—one that relies on an understanding of the interactions of different narrative

features—will allow researchers to develop a more multi-dimensional understanding of narrative response beyond simple binaries like fact-based versus narrative discourse. Doing so will allow research to address the kinds of narratives that readers, viewers and listeners respond most strongly to and better understand the textual features that drive responses. Given the proliferation of user generated online responses "in the wild" today, a rich research opportunity awaits to move beyond laboratory environments (Bail, 2014; Salganik, 2019) to understand the relationship between different narrative products (movies, books, advertisements) and consumer responses captured by social media.

3.3 Narrative beliefs

Common to the work discussed so far is a reliance on two guiding assumptions: that narrativity is both explicit (encoded within given textual features) and discrete (bounded by a single document or series of documents). Yet, as research in narrative psychology has long argued, narrative can also be conceived as a latent organizing principle of human reasoning (Bruner, 1991).

Sociologists refer to the notion of "ontological narratives" (Somers, 1994) or "deep stories" (Hochschild, 2018) to refer to the narrative structuring of information that individuals use to understand their place in the world and facilitate decision-making. Such deep stories are not located exclusively or even explicitly within a single narrative instance; rather, they can be thought of as higher-order organizing principles that situate multiple events of speech into a more general paradigm. One powerful example of this is the so-called "Great Replacement" narrative, the worldview held by some white Americans or Europeans that non-white immigrants are arriving to take their jobs and political power (Feola, 2020). While this belief takes many different forms, it can be distilled into a single and relatively cogent story.

A greater understanding of these deep stories and how they shape communities or political behavior can have significant social outcomes. They can help policymakers craft policies that better reflect community belief systems and thus produce higher levels of adherence and support as well as contribute to efforts at mitigating social crises. For example, Shiller (2020) argues that the severity of economic downturns might in part rest on the vividness of certain stories circulating in society,

and thus, one effective strategy in containing this severity might consist in promoting alternative narratives to the public. Researchers are increasingly aware that changing human behavior depends on changing the stories people tell themselves and their social circles.

4 Conclusion

Narrative is a pervasive phenomenon that cuts across the human experience. As research in NLP has increasingly focused on developing computational models to reason about stories, we see space for existing structures in narrative theory to help provide an organizing system around these efforts. By making connections between theoretical frameworks in narratology and existing research in NLP more explicit, we hope to provide a mechanism to coordinate diverse research efforts, unify terminology and methods, and suggest major research challenges that will help drive the empirical study of narrative forward.

5 Acknowledgments

We thank the anonymous reviewers for their helpful comments. The research reported in this article was supported by funding from the Social Sciences and Humanities Research Council of Canada (895-2013-1011), the National Science Foundation (IIS-1942591), and the National Endowment for the Humanities (HAA-271654-20).

References

- Jonathan M Adler, Jennifer Lodi-Smith, Frederick L Philippe, and Iliane Houle. 2016. The incremental validity of narrative identity in predicting well-being: A review of the field and recommendations for the future. *Personality and Social Psychology Review*, 20(2):142–175.
- David Ahn. 2006. [The stages of event extraction](#). In *Proceedings of the Workshop on Annotating and Reasoning About Time and Events*, ARTE '06, pages 1–8, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Christopher A Bail. 2014. The cultural environment: Measuring culture with big data. *Theory and Society*, 43(3-4):465–482.
- Mieke Bal and Christine Van Boheemen. 2009. *Narratology: Introduction to the theory of narrative*. University of Toronto Press.
- David Bamman, Olivia Lewke, and Anya Mansoor. 2020. [An annotated dataset of coreference in English literature](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 44–54, Marseille, France. European Language Resources Association.
- David Bamman, Sejal Papat, and Sheng Shen. 2019. [An annotated dataset of literary entities](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2138–2144, Minneapolis, Minnesota. Association for Computational Linguistics.
- Roland Barthes and Lionel Duisit. 1975. An introduction to the structural analysis of narrative. *New literary history*, 6(2):237–272.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Ryan L Boyd, Kate G Blackburn, and James W Pennebaker. 2020. The narrative arc: Revealing core narrative structures through text analysis. *Science advances*, 6(32):eaba2196.
- Amber E. Boydston. 2013. *Making the News: Politics, the Media, and Agenda Setting*. University of Chicago Press.
- Kurt Braddock and James Price Dillard. 2016. Meta-analytic evidence for the persuasive effect of narratives on beliefs, attitudes, intentions, and behaviors. *Communication Monographs*, 83(4):446–467.
- Faeze Brahman and Snigdha Chaturvedi. 2020. [Modeling protagonist emotions for emotion-aware storytelling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5277–5294, Online. Association for Computational Linguistics.
- Peter Broadwell, David Mimno, and Timothy Tangherlini. 2018. The tell-tale hat: Surfacing the uncertainty in folklore classification. *Journal of Cultural Analytics*.
- Julian Brooke, Adam Hammond, and Timothy Baldwin. 2016. [Bootstrapped text-level named entity recognition for literature](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 344–350, Berlin, Germany. Association for Computational Linguistics.
- Jerome Bruner. 1991. The narrative construction of reality. *Critical inquiry*, 18(1):1–21.

- Annellen Brunner, Ngoc Duyen Tanja Tu, Lukas Weimer, and Fotis Jannidis. 2020. To BERT or not to BERT—comparing contextual embeddings in a deep learning architecture for the automatic recognition of four types of speech, thought and writing representation. In *Proceedings of the 5th Swiss Text Analytics Conference (SwissText) & 16th Conference on Natural Language Processing (KONVENS)*.
- Kenneth Burke. 1969. *A grammar of motives*, volume 177. Univ of California Press.
- Simon Bushell, Géraldine Satre Buisson, Mark Workman, and Thomas Colley. 2017. [Strategic narratives in climate change: Towards a unifying narrative to address the action gap on climate change](#). *Energy Research Social Science*, 28:39–49.
- Dallas Card, Amber E. Boydstun, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. [The media frames corpus: Annotations of frames across issues](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444. Association for Computational Linguistics.
- Dallas Card, Justin Gross, Amber Boydstun, and Noah A. Smith. 2016. Analyzing framing through the casts of characters in the news. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1410–1420, Austin, Texas. Association for Computational Linguistics.
- Nathanael Chambers. 2013. [Event schema induction with a probabilistic entity-driven model](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Seattle, Washington, USA. Association for Computational Linguistics.
- Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.
- Nathanael Chambers and Dan Jurafsky. 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*, ACL ’09, pages 602–610, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Snigdha Chaturvedi, Dan Goldwasser, and Hal Daume III. 2016a. Ask, and shall you receive? understanding desire fulfillment in natural language text. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Snigdha Chaturvedi, Shashank Srivastava, Hal Daume III, and Chris Dyer. 2016b. Modeling evolving relationships between characters in literary novels. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. [Event extraction via dynamic multi-pooling convolutional neural networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 167–176. Association for Computational Linguistics.
- Jonathan Y Cheng. 2020. Fleshing out models of gender in English-language novels (1850–2000). *Journal of Cultural Analytics*.
- Jackie Chi Kit Cheung, Hoifung Poon, and Lucy Vanderwende. 2013. [Probabilistic frame induction](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 837–846, Atlanta, Georgia. Association for Computational Linguistics.
- Mariona Coll Ardanuy, Federico Nanni, Kaspar Beelen, Kasma Hosseini, Ruth Ahnert, Jon Lawrence, Katherine McDonough, Giorgia Tolfo, Daniel CS Wilson, and Barbara McGillivray. 2020. [Living machines: A study of atypical animacy](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4534–4545, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Sonya Dal Cin, Mark P Zanna, and Geoffrey T Fong. 2004. Narrative persuasion and overcoming resistance. *Resistance and persuasion*, 2(175-191):4.
- Rodolfo Delmonte and Giulia Marchesini. 2017. A semantically-based computational approach to narrative structure. In *IWCS 2017—12th International Conference on Computational Semantics—Short papers*.
- Grant DeLozier, Ben Wing, Jason Baldridge, and Scott Nesbit. 2016. [Creating a novel geolocation corpus from historical texts](#). In *Proceedings of the 10th Linguistic Annotation Workshop held in conjunction with ACL 2016 (LAW-X 2016)*, pages 188–198. Association for Computational Linguistics.
- Richard Doust and Paul Piwek. 2017. A model of suspense for narrative generation. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 178–187.
- Jens Eder, Fotis Jannidis, and Ralf Schneider. 2010. *Characters in fictional worlds: Understanding imaginary beings in literature, film, and other media*, volume 3. Walter de Gruyter.
- Joshua Eisenberg and Mark Finlayson. 2016. [Automatic identification of narrative diegesis and point of view](#). In *Proceedings of the 2nd Workshop on Computing News Storylines (CNS 2016)*, pages 36–46, Austin, Texas. Association for Computational Linguistics.

- David Elson, Nicholas Dames, and Kathleen McKeown. 2010. [Extracting social networks from literary fiction](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 138–147, Uppsala, Sweden. Association for Computational Linguistics.
- Matt Erlin, Andrew Piper, Douglas Knox, Stephen Pentecost, Michaela Drouillard, Brian Powell, and Cienna Townson. 2021. Cultural capitals: Modeling ‘minor’ european literature. *Cultural Analytics*.
- Xiaocheng Feng, Lifu Huang, Duyu Tang, Heng Ji, Bing Qin, and Ting Liu. 2016. [A language-independent neural network for event detection](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 66–71. Association for Computational Linguistics.
- Michael Feola. 2020. “You will not replace us”: The melancholic nationalism of whiteness. *Political Theory*, page 0090591720972745.
- Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. [Framing and agenda-setting in Russian news: a computational analysis of intricate political strategies](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3570–3580, Brussels, Belgium. Association for Computational Linguistics.
- Monika Fludernik. 2002. *Towards a ‘Natural’ Narratology*. Routledge.
- Edward Morgan Forster. 1985. *Aspects of the Novel*, volume 19. Houghton Mifflin Harcourt.
- Gustav Freytag. 1895. *Technique of the drama: An exposition of dramatic composition and art*. S. Griggs.
- John Frow. 2014. *Character and person*. Oxford University Press.
- Gärard Genette. 1997. *Palimpsests: Literature in the second degree*, volume 8. U of Nebraska Press.
- Gérard Genette. 1983. *Narrative discourse: An essay in method*, volume 3. Cornell University Press.
- Alexandra Georgakopoulou. 2007. *Small stories, interaction and identities*, volume 8. John Benjamins Publishing.
- Rachel Giora and Yeshayahu Shen. 1994. Degrees of narrativity and strategies of semantic reduction. *Poetics*, 22(6):447–458.
- René Girard. 2020. Generative scapegoating. In *Violent origins*, pages 73–105. Stanford University Press.
- Arthur C Graesser, Brent Olde, and Bianca Klettke. 2002. How does the mind construct and represent stories. *Narrative impact: Social and cognitive foundations*, pages 229–262.
- Michael Hauge. 2017. *Storytelling Made Easy: Persuade and Transform Your Audiences, Buyers, and Clients—Simply, Quickly, and Profitably*. Indie Books International.
- Hua He, Denilson Barbosa, and Grzegorz Kondrak. 2013. Identification of speakers in novels. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1312–1320, Sofia, Bulgaria. Association for Computational Linguistics.
- David Herman. 2009. *Basic elements of narrative*. John Wiley & Sons.
- Arlie Russell Hochschild. 2018. *Strangers in their own land: Anger and mourning on the American right*. The New Press.
- Axel Honneth. 1996. *The struggle for recognition: The moral grammar of social conflicts*. MIT Press.
- Peter Hühn. 2014. *Event and eventfulness*. de Gruyter.
- Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. 2016. [Feuding families and former Friends: Unsupervised learning for dynamic fictional relationships](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1534–1544, San Diego, California. Association for Computational Linguistics.
- Labiba Jahan, Geeticka Chauhan, and Mark Finlayson. 2018. [A new approach to Animacy detection](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1–12, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Labiba Jahan and Mark Finlayson. 2019. [Character identification refined: A proposal](#). In *Proceedings of the First Workshop on Narrative Understanding*, pages 12–18, Minneapolis, Minnesota. Association for Computational Linguistics.
- Fotis Jannidis. 2009. *Character*. De Gruyter.
- Matthew Jockers. 2015. Revealing sentiment and plot arcs with the syuzhet package. <http://www.matthewjockers.net/2015/02/02/syuzhet/>.
- Folger B Karsdorp, M Meulen, Theo Meder, and APJ van den Bosch. 2015. Animacy detection in stories.
- Edward Kearns. 2020. [Annotating and quantifying narrative time disruptions in modernist and hypertext fiction](#). In *Proceedings of the First Joint Workshop on Narrative Understanding, Storylines, and Events*, pages 72–77, Online. Association for Computational Linguistics.

- Allen Kim, Charuta Pethe, and Steve Skiena. 2020. [What time is it? temporal analysis of novels](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9076–9086, Online. Association for Computational Linguistics.
- Eve Kraicer and Andrew Piper. 2019. Social characters: the hierarchy of gender in contemporary english-language fiction. *Journal of Cultural Analytics*.
- Emily Kubin, Curtis Puryear, Chelsea Schein, and Kurt Gray. 2021. [Personal experiences bridge moral and political divides better than facts](#). *Proceedings of the National Academy of Sciences*, 118(6):e2008389118.
- Vincent Labatut and Xavier Bost. 2019. [Extraction and analysis of fictional character networks: A survey](#). *CoRR*, abs/1907.02704.
- W. Labov and J. Waletzky. 1967. Narrative analysis: Oral versions of personal experience. In *Essays on the verbal and visual arts*, Seattle. University of Washington Press.
- Henri Lefebvre and Donald Nicholson-Smith. 1991. *The production of space*, volume 142. Oxford Blackwell.
- Effi Levi, Guy Mor, Shaul Shenhav, and Tamir Sheafer. 2020. [Compres: A dataset for narrative structure in news](#).
- Qi Li, Heng Ji, and Liang Huang. 2013. [Joint event extraction via structured prediction with global features](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 73–82. Association for Computational Linguistics.
- Genevieve Liveley. 2019. *Narratology*. Oxford University Press.
- Hoyt Long and Richard Jean So. 2016. Turbulent flow: A computational model of world literature. *Modern Language Quarterly*, 77(3):345–367.
- Aibek Makazhanov, Denilson Barbosa, and Grzegorz Kondrak. 2014. [Extracting family relationship networks from novels](#). *CoRR*, abs/1405.0603.
- Philip Massey, Patrick Xia, David Bamman, and Noah A. Smith. 2015. [Annotating character relationships in literary texts](#).
- Elena Mikhalkova, Timofei Protasov, Polina Sokolova, Anastasiia Bashmakova, and Anastasiia Drozdova. 2020. Modelling narrative elements in a short story: A study on annotation schemes and guidelines. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 126–132.
- Franco Moretti. 2013. Operationalizing. *New Left Review*, (84).
- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016. [CaTeRS: Causal and temporal relation scheme for semantic annotation of event structures](#). In *Proceedings of the Fourth Workshop on Events*, pages 51–61, San Diego, California. Association for Computational Linguistics.
- Grace Muzny, Mark Algee-Hewitt, and Dan Jurafsky. 2017a. Dialogism in the novel: A computational model of the dialogic nature of narration and quotations. *Digital Scholarship in the Humanities*, 32.
- Grace Muzny, Michael Fang, Angel Chang, and Dan Jurafsky. 2017b. A two-stage sieve approach for quote attribution. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 460–470.
- Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. [Joint event extraction via recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 300–309. Association for Computational Linguistics.
- Thien Huu Nguyen and Ralph Grishman. 2015. [Event detection and domain adaptation with convolutional neural networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 365–371. Association for Computational Linguistics.
- Jessica Ouyang and Kathy McKeown. 2014. [Towards automatic detection of narrative structure](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 4624–4631, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Ruth Page. 2010. Re-examining narrativity: Small stories in status updates. *Text & Talk*, 30(4):423–444.
- Alan Palmer. 2004. *Fictional minds*. U of Nebraska Press.
- Pinelopi Papalampidi, Frank Keller, and Mirella Lapata. 2019. [Movie plot analysis via turning point identification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1707–1717, Hong Kong, China. Association for Computational Linguistics.
- Federico Ponzola. 2018. Looking at narrative as a complex system: The proteus principle. In *Narrating complexity*, pages 101–122. Springer.
- Barbara Piatti and Lorenz Hurni. 2009. Mapping the ontologically unreal–counterfactual spaces in literature and cartography. *The Cartographic Journal*, 46(4):333–342.

- Barbara Piatti, Anne-Kathrin Reuschel, and Lorenz Hurni. 2013. Dreams, longings, memories—visualising the dimension of projected spaces in fiction. In *Proceedings of the 26th International Cartographic Conference*, pages 74–92.
- Andrew Piper. 2015. Novel devotions: Conversional reading, computational modeling, and the modern novel. *New Literary History*, 46(1):63–98.
- Andrew Piper. 2018a. 5. Characterization (constraint). In *Enumerations*, pages 118–146. University of Chicago Press.
- Andrew Piper. 2018b. *Enumerations: data and literary study*. University of Chicago Press.
- Andrew Piper and Eva Portelance. 2016. How cultural capital works: Prizewinning novels, bestsellers, and the time of reading. *Post45*.
- Vladimir Propp. 2010. *Morphology of the Folktale*, volume 9. University of Texas Press.
- Hannah Rashkin, Antoine Bosselut, Maarten Sap, Kevin Knight, and Yejin Choi. 2018. [Modeling naive psychology of characters in simple common-sense stories](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2289–2299, Melbourne, Australia. Association for Computational Linguistics.
- Andrew J Reagan, Lewis Mitchell, Dilan Kiley, Christopher M Danforth, and Peter Sheridan Dodds. 2016. The emotional arcs of stories are dominated by six basic shapes. *EPJ Data Science*, 5(1):1–12.
- Nils Reiter, Marcus Willand, and Evelyn Gius. 2019. A shared task for the digital humanities chapter 1: Introduction to annotation, narrative levels and shared tasks. *Journal of Cultural Analytics*.
- Paul Ricoeur. 2014. *Time and narrative, Volume 3*, volume 3. University of Chicago Press.
- Kenji Sagae, Andrew S Gordon, Morteza Dehghani, Mike Metke, Jackie S Kim, Sarah I Gimbel, Christine Tipper, Jonas Kaplan, and Mary Helen Immordino-Yang. 2013. A data-driven approach for classification of subjectivity in personal narratives. In *2013 Workshop on Computational Models of Narrative*.
- Belen Saldias and Deb Roy. 2020. [Exploring aspects of similarity between spoken personal narratives by disentangling them into narrative clause types](#). In *Proceedings of the First Joint Workshop on Narrative Understanding, Storylines, and Events*, pages 78–86, Online. Association for Computational Linguistics.
- Matthew J Salganik. 2019. *Bit by bit: Social research in the digital age*. Princeton University Press.
- Maarten Sap, Marcella Cindy Prasettio, Ari Holtzman, Hannah Rashkin, and Yejin Choi. 2017. [Connotation frames of power and agency in modern films](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2329–2334, Copenhagen, Denmark. Association for Computational Linguistics.
- Roger C. Schank and Robert P. Abelson. 1977. *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. L. Erlbaum, Hillsdale, NJ.
- Amartya Sen. 1993. Capability and well-being. *The quality of life*, 30:270–293.
- Robert J Shiller. 2020. *Narrative economics: How stories go viral and drive major economic events*. Princeton University Press.
- Matthew Sims, Jong Ho Park, and David Bamman. 2019. [Literary event detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3623–3634, Florence, Italy. Association for Computational Linguistics.
- Margaret R Somers. 1994. The narrative constitution of identity: A relational and network approach. *Theory and society*, pages 605–649.
- Meir Sternberg. 1992. Telling in time (ii): Chronology, teleology, narrativity. *Poetics today*, 13(3):463–541.
- Reid Swanson, Elahe Rahimtoroghi, Thomas Corcoran, and Marilyn Walker. 2014. [Identifying narrative clause types in personal stories](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 171–180, Philadelphia, PA, U.S.A. Association for Computational Linguistics.
- Rochelle Terman. 2017. Islamophobia and media portrayals of Muslim women: A computational text analysis of US news coverage. *International Studies Quarterly*, 61(3):489–502.
- Stith Thompson. 1989. *Motif-index of Folk-literature: JK*, volume 4. Indiana University Press.
- Tzvetan Todorov. 1981. *Introduction to poetics*, volume 1. U of Minnesota Press.
- Shubham Toshniwal, Sam Wiseman, Allyson Ettinger, Karen Livescu, and Kevin Gimpel. 2020. [Learning to Ignore: Long Document Coreference with Bounded Memory Neural Networks](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8519–8526, Online. Association for Computational Linguistics.
- Ted Underwood. 2018. Why literary time is measured in minutes. *ELH*, 85(2).
- Ted Underwood, David Bamman, and Sabrina Lee. 2018. The transformation of gender in english-language fiction. *Journal of Cultural Analytics*.

- Naushad UzZaman, Hector Llorens, Leon Derczynski, James Allen, Marc Verhagen, and James Pustejovsky. 2013. [SemEval-2013 task 1: TempEval-3: Evaluating time expressions, events, and temporal relations](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 1–9, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Hardik Vala, David Jurgens, Andrew Piper, and Derek Ruths. 2015. [Mr. Bennet, his coachman, and the Archbishop walk into a bar but only one of them gets recognized: On the difficulty of detecting characters in literary texts](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 769–774, Lisbon, Portugal. Association for Computational Linguistics.
- Teun A Van Dijk. 1976. Philosophy of action and theory of narrative. *Poetics*, 5(4):287–338.
- Marc Verhagen, Robert Gaizauskas, Frank Schilder, Mark Hepple, Jessica Moszkowicz, and James Pustejovsky. 2009. The TempEval challenge: identifying temporal relations in text. *Language Resources and Evaluation*, 43(2):161–179.
- Krishnapriya Vishnubhotla, Adam Hammond, and Graeme Hirst. 2019. [Are fictional voices distinguishable? classifying character voices in modern drama](#). In *Proceedings of the 3rd Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 29–34, Minneapolis, USA. Association for Computational Linguistics.
- Byron C Wallace. 2012. Multiple narrative disentanglement: Unraveling Infinite Jest. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1–10.
- Richard Walsh. 2018. Narrative theory for complexity scientists. In *Narrating complexity*, pages 11–25. Springer.
- Matthew Wilkens. 2013. The geographic imagination of Civil War-era American fiction. *American Literary History*, 25(4):803–840.
- David Wilmot and Frank Keller. 2020. [Modelling suspense in short stories as uncertainty reduction over neural representation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1763–1788, Online. Association for Computational Linguistics.
- Alex Woloch. 2009. *The One vs. the Many: Minor Characters and the Space of the Protagonist in the Novel*. Princeton University Press.
- Bishan Yang and Tom M. Mitchell. 2016. [Joint extraction of events and entities within a document context](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 289–299. Association for Computational Linguistics.
- Gregory Yauney, Ted Underwood, and David Mimno. 2019. Computational prediction of elapsed narrative time. *Workshop on Narrative Understanding*.