Computational Narrative Understanding: A Big Picture Analysis

Andrew Piper
McGill University
680 Sherbrooke St. West
Montreal, QC H3A 2M7 CANADA
andrew.piper@mcgill.ca

Abstract

This paper provides an overview of outstanding major research goals for the field of computational narrative understanding. Storytelling is an essential human practice, one that provides a sense of personal meaning, shared sense of community, and individual enjoyment. A number of research domains have increasingly focused on storytelling as a key mechanism for explaining human behavior. Now is an opportune moment to provide a vision of the contributions that computational narrative understanding can make towards this collective endeavor and the challenges facing the field. In addition to providing an overview of the elements of narrative, this paper outlines three major lines of inquiry: understanding the multimodality of narrative; the temporal patterning of narrative (narrative "shape"); and sociocultural narrative schemas, i.e. collective narratives. The paper concludes with a call for more inter-disciplinary working groups and deeper investment in building cross-cultural and multimodal narrative datasets.

1 Introduction

The Native-American writer, Gerald Vizenor, once remarked: "There isn't any center to the world but a story" (Coltelli, 1990). Storytelling is a ubiquitous human practice, exhibited in all human cultures, languages, and recorded historical time periods. Many of the world's most enduring and widespread belief systems are encoded through stories, and research suggests that human reasoning (Bruner, 1991) and selfhood (Berns, 2022) are fundamentally grounded in narrative. Today, a growing body of research is developing across a variety of domains that focus on storytelling as a key mechanism for explaining human beliefs and behavior, from mental health (Adler et al., 2016), to political stance taking (Bushell et al., 2017), to consumer persuasion (Bilandzic and Busselle, 2013), to financial decision making (Shiller, 2020).

Given this widespread interest in, and awareness of, narrative as a crucial driver of human behavior, the field of "computational narrative understanding" has a great opportunity to contribute to a range of research fields. Computational narrative understanding has crystallized over the past 5-10 years as a vibrant subset of natural language processing (Bamman et al., 2019; Jorge et al., 2019). Its aim is to develop computational systems for the detection and understanding of narrative communication across different media and different cultural domains. While we may typically think of stories as encoded in written documents, the practice of narrative can be represented through a diverse array of media, including oral speech, song, still or moving images, social media, playable media like video games, or some combination of the above.

The aim of this paper is to provide a big picture view of some of the key higher-level goals for computational narrative understanding. A great deal of on-going and inspiring work continues to make progress in the detection and analysis of different components of narrative communication (for a review see Piper et al. (2021)). It thus seems timely to provide a vision of where we are going as a community to help motivate and organize future work in the field.

In section two, I provide a brief minimal definition of narrative communication highlighting its constituent parts building on prior work (Piper et al., 2021). Before moving to the big picture, it is important to ground our understanding of this core concept. In section three, I describe a research framework that aims to develop a more multi-modal understanding of narrative. With its grounding in NLP, computational narrative understanding has understandably focused on narrative as a linguistic phenomenon. However, as narratologists have long pointed out (Ryan et al., 2004), storytelling can transpire in numerous different media. Being able to integrate observations across

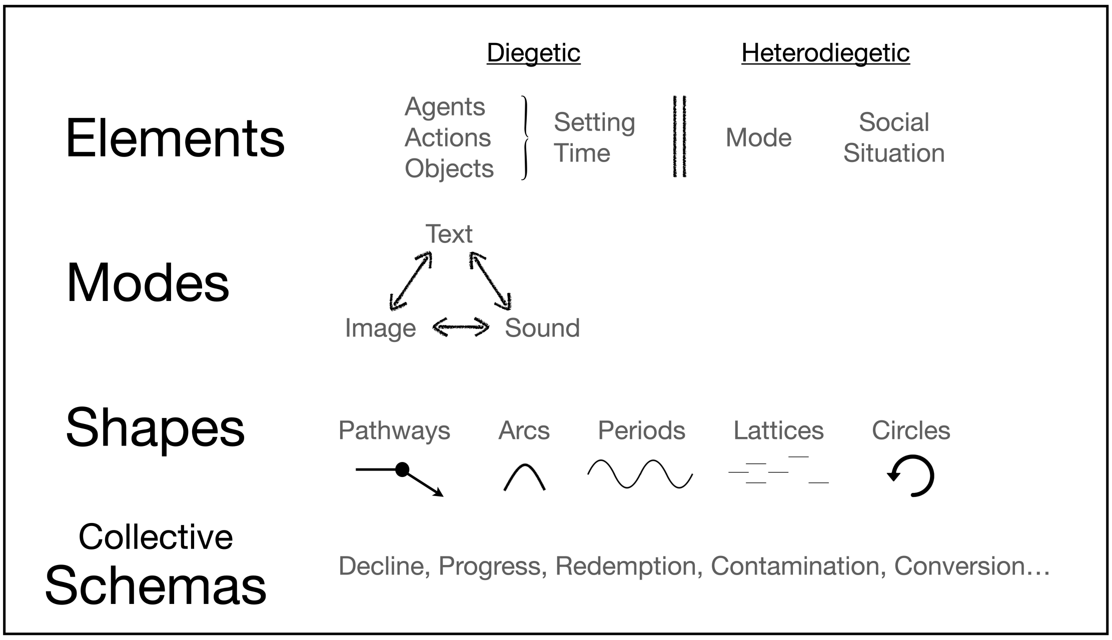

Figure 1: Overview of narrative research areas discussed in this paper.

media, from speech to text to images to playable media should become a central goal of computational narrative understanding.

In section four, I describe a research framework aimed at understanding narrative "shape" (also called "form" or "structure" (Berhe et al., 2022)), which can be understood as the temporal patterning of narrative elements. One of the fundamental aspects of storytelling is the encoding of events in time (Genette, 1983; Sternberg, 1992; Ricoeur, 2012). Narrative meaning is thus contingent on the temporal organization of information.

Seeing narratives as temporal artifacts, made in time and composed of time, then leads to the highest-level form of integration described in section five, that of narrative "schemas." As Berns (2022) has argued, narratives are forms of information compression, reducing the vast scope of experienced data down to a much more limited set of communicated data. Such compression necessarily follows archetypes or patterns that can be biologically or culturally conditioned (or some mixture of the two).

While the idea of "scripts" has been applied to understand the local schematic encoding of events (Chambers and Jurafsky, 2008a), prior work in folklore studies has offered promising frameworks for expanding the idea of schema to include whole stories within various typologies (Thompson, 1989). Essential to this framework is an attention to larger narrative ecologies, the ways in which such schemas play a generative and/or organizing role within broader, and potentially interactive, communicative domains (Tangherlini et al., 2020).

It is common to think of narrative as located within an individual document or artifact (this book or blog post tells a story), but narratologists have also highlighted the way story structures emerge from the complex social interactions of numerous agents (known as the "small stories" paradigm (Georgakopoulou, 2007)). Such "small stories" can then coalesce into larger socially circulatable schemas, variously referred to as "ontological narratives" (Somers, 1994), "deep stories" (Hochschild, 2018), or "collective narratives" (Bliuc and Chidley, 2022). Such schemas can then guide the processing and circulation of new information to "fit the narrative," potentially creating informational feedback loops that are durable over shorter or longer stretches of time.

In sum, we want to have a research framework capable of scaling the ladder from local elements (section 2), different media (section 3), formal structure (section 4), all the way up to schemas and social dynamics (section 5). Figure 1 provides a schematic overview of this big picture.

I conclude in section six with a reflection on the need for greater inter-disciplinary collaboration and deeper investment in building cross-cultural and multi-modal datasets. As we develop more sophisticated systems for detecting narrative communication, we will want to invest more deeply in the infrastructure for large-scale narrative understanding. This will necessarily entail collaborations across disciplines to better understand socially relevant applications as well as the ability to develop appropriate data. It will also require developing an awareness around the limitations or risks of narrative communication (Salmon, 2017; Gottschall, 2021). Stories not only inspire and move audiences, they can also deform reality and misinform, a point that should remain at the forefront of our thinking about how stories stand at the centre of so much human behavior, for better and for worse.

2 The Elements of Narrative

At its most elementary level, a story can be said to occur when all of the following criteria are met:

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.

For there to be a story, we need (A) a teller, (B) a mode of telling (i.e. medium), (C) a recipient, (D) a social situation, (E) an agent, (F) at least one action or event, (G) a possible object, (H) a location, (I) a time-frame, and (J) a motivation or cause of the actions involved. Narratologists make a distinction between the frame of the storyworld (i.e. all of the elements that come after the double lines above) known as "diegetic" elements, and the frame of telling (i.e. all of the elements that come before the double lines) known as "heterodiegetic" elements, where diegesis refers to a narrative "frame" or "world."

Importantly, not all of these elements need to be explicit. For example, in one of the most famous short stories ever proposed by Ernest Hemingway, very little from the above list is specified:

For sale: Baby shoes. Never worn.

We don't know where and when this happened, nor do we know who is telling the story. All we know is what happened (on two levels): a baby died and a family needs money. But no matter how much is implicit in this story all of the parts are there. Something happens to someone somewhere at some time for some reason and someone tells someone this story.

Such a definition can be useful because it highlights the array of narrative elements that require computational solutions to "understand" the cultural meaning of a story. Such applications have included: character detection (Bamman et al., 2014; Jahan et al., 2018; Piper, 2023b; Stammbach et al., 2022), object detection (Piper and Bagga, 2022a), character relation detection (Labatut and Bost, 2019; Kraicer and Piper, 2019), event detection (Vauth et al., 2021), geographic and spatial understanding (Wilkens, 2013; Evans and Wilkens, 2018; Piatti et al., 2013; Erlin et al., 2021), temporal understanding (Underwood, 2018; Yauney et al., 2019; Vossen et al., 2021; Gangal et al., 2022), and causality mining (Meehan and Piper, 2022). A full review can be found in Piper et al. (2021) and Santana et al. (2023).

A second, higher-level way that a story can be broken down into constituent parts is through *discourse elements*. As we will see, this problem is associated with challenges of text segmentation, though importantly differs from prior work focused on sequential and/or paratextual (i.e. chapter) segmentation (Pethe et al., 2020; Zehe et al., 2021).

Narratives not only contain event-frames (i.e. scenes), but are also composed of heterogeneous linguistic styles in which the act of narration is but one component. This is one reason recent narrative theory has emphasized the idea of "narrativity" (Piper and Bagga, 2022b; Pianzola, 2018; Giora and Shen, 1994), which captures the *degree* of narration intrinsic to a narrative. An ostensibly narrative document like a short story will engage in moments of non-narrative statements, just as putatatively non-narrative documents like scientific articles may engage in occasionaly moments of narration. Narration is in this sense not a universal property of documents, but a local linguistic phenomenon. As Ochs et al. (2009) write, "We believe that narrative as genre and activity can be fruitfully examined in terms of a set of dimensions that a narrative displays to different degrees and in different ways."

Narratologists typically break down narratives into at least four basic discourse components:

Discourse	**Contents**
1. Narration	Agents and events
2. Description	Setting, modification, context
3. Dialogue	Reported speech
4. Evaluation	Meta-level discourse

"Narration," also known as "diegesis," refers to the linguistic structures described above that occur after the double horizontal line (E-J). This is the classic understanding of narrative, where events pertaining to an agent are recounted (this can also fall under the heading of "eventfulness" (Hühn, 2014)).

"Description," also called "mimesis," refers to when the surroundings or context of events are described and during which events do not unfold (though they may be unfolding in the background). In cinema, this is equivalent to an "establishing shot" that indicates to viewers where they are in time and space. Crucial to description is that it lacks the agent/action/cause structure from above.

"Dialogue" refers to any form of reported speech, though it can also take the form of indirect speech as well. Recounting what characters say to each other is an integral component of stories, although it technically is a form of dramatic performance (for a reflection on this topic see (Genette, 1992)).

Finally, many stories contain what we might call meta-textual statements (called "evaluation" by Labov and Waletzky (1967)), where the narrator provides some higher-level assessment with regards to the story, either a reflection on the story contents, their meaning, or some didactic lesson that should be imparted, making a latent feature of storytelling (it's meaning or purpose) manifest. While it may come at the end of a story, it can also be interspersed throughout. Here are a few examples of such statements:

1. *It is a truth universally acknowledged, that a single man in possession of a good fortune, must be in want of a wife.* (Pride and Prejudice)

2. *The flatterer lives at the expense of those who will listen to him.* (Aesop's Fables)

3. *All in all, I'd say that those years were some of the best times I've ever had.* (AskReddit)

While there are many more ways one can parse a story (see Bal and Van Boheemen (2009); Genette (1983)), the frameworks above provide practical heuristics for the ways that stories can be broken down into more elementary parts to ground computational models.

3 The Modality of Narrative

Grounded in NLP, computational narrative understanding has largely prioritized written narratives for understandable reasons. However, such text-driven approaches leave out large portions of storytelling behavior, including movies and television (Arnold et al., 2019; Papalampidi et al., 2019), user-generated streaming content, illustrated content in comic strips (Edlin and Reiss, 2023), graphic novels, or children's books (Adukia et al., 2021), and finally video games, which might have stronger or weaker narrative structures. While textual narratives are largely unimodal in nature (though the physical and visual dimensions of books has been a vibrant area of study for a long time (Collective, 2019)), these other narrative forms are all crucially multi-modal in nature.

Sound, image, and language can interact in ways that are complex and dynamic. A robust field of multimodal NLP research into text-image interactions for meaning-making has emerged in recent years (see for example recent research on humour by Hasan et al. (2019); Hessel et al. (2023)). Nevertheless, investigations into multimodal *narrative* understanding, such as the relationship between text and illustrations in children's books or graphic novels is in need of more attention (see Adukia et al. (2021) for an example exploring the visual qualities of children's book illustrations with respect to race). Understanding the kinds of gestural or pictorial preferences that are foregrounded given certain textual cues could give us insights into the way humans translate language into image (and vice versa) across different cultural domains.

Similarly, we still lack major comparative studies of narrative behavior across media, i.e. comparisons of narrative elements and archetypes in film, television, user-generated content, oral performances and books. For example, evidence suggests that written and oral narratives have similar "establishing shot" structures similar to movies and television (Boyd et al., 2020; Piper, 2023a). More precise comparisons can highlight the modal-specificity of different narrative elements along with the transmodal practices that are independent of a given modality. Understanding the ways in which storytellers marshal images, sounds, and words to create immersive experiences for audiences will greatly contribute to the project of computational narrative understanding.

4 The Shape of Narrative

The writer and critic Italo Calvino was fond of quoting a Sicilian expression that "time takes no time in a story" (Calvino, 1988). A narrator can tell a story that traverses centuries in a few sentences or can slow time down to the point where a few seconds takes minutes to describe. Narratologists refer to this as the difference between *narrated time* (the time transpiring in the storyworld) and *narrative time* (how long a story takes to tell). No matter how much stories may compress time, they cannot be told all at once. Contrary to Calvino's favored Sicilian expression, all stories, even the shortest, take time to tell.

This temporal dimension of narrative – that stories take time to tell and tell of things happening in time – has long been one of the privileged topics of narrative theory (Ricoeur, 2012; Sternberg, 1990, 1992). As the theorist David Herman has argued, "Narrative is a basic human strategy for coming to terms with time, process, and change" (Herman, 2009).

A number of approaches have been proposed for the computational modeling of temporal patterns in narrative (for a review of modeling narrative structure see Berhe et al., 2022). Schmidt (2015) used topic modeling to identify thematic arcs in television screenplays, while Thompson et al. (2018) used topic models to study thematic progression in philosophical texts and social media. Reagan et al. (2016) used sentiment analysis to model the concept of narrative fortune (Freytag, 1895), for which Elkins (2022) provides a more in-depth study of the validity of sentiment arcs as models of narrative structure. Boyd et al. (2020) used particular word types to capture three primary narrative stages, and Sap et al. (2022) used the predictability of next sentences to capture the concept of narrative "flow."

Piper and Toubia (2023) used word embeddings to model narrative non-linearity using the traveling salesman problem, while Toubia et al. (2021) offer two further ways of thinking about narrative shape called "speed" and "volume." Researchers have also used information theoretic frameworks to model the concept of narrative revelation using time series methods (Piper, 2023a) and stylistic novelty over narrative time using a bloom filter (McGrath, 2018). Ouyang and McKeown (2015) and Piper (2015) devised methods for predicting narrative "turning points" as larger structural qualites, drawing on Aristotelian and Augustinian theories of narrative respectively.

Common to all of these models is the assumption that the dissemination of information over narrative time assumes observable patterns (called "form" or "structure") and that these patterns encode cultural meaning. The most common framework to date has been that of the narrative "arc," drawn from French neo-classical tragedy (Freytag, 1895). According to this model, narratives encode a central conflict that results in some form of resolution or change, which can be approximated by an arc of rising and falling fortune or conflict.

Much future work remains to better understand relevant ways of capturing narrative time in terms of its formal patterns. The first area of consideration should be further work into the choice of feature distributions that are used to capture narrative time. Where prior work has focused to date on topic models, sentiment vocabulary, word embeddings, lexemes, and letters, higher-level narrative features (see Section 1) should continue to be developed and studied. We assume that the distribution of characters, event types, locations, or narrative modes may also contribute to the overall structural qualities of stories.

Second, modeling narrative change itself remains a key area of further research. Prior empirical work has shown that long narratives may employ multiple "arcs" rather than single turning points (Reagan et al., 2016; Fudolig et al., 2023), while other work has emphasized the significance of single turning points (Ouyang and McKeown, 2015; Piper, 2015). Additionally, the identity or meaning of such moments of change, regardless of how many, are also not well understood. The dramatic model of narrative denouement suggests that turning points are best understood as forms of "conflict/resolution," while other narrative theories suggest that "surprisingness" is the optimal way of understanding narrative change (Wilmot and Keller, 2020). Brewer and Lichtenstein (1982) have proposed two further affective states of *suspense* and *curiosity* in addition to surprise to capture the discrepancy between storyworld information and narrative information (i.e. when key information is withheld or forms of temporal anachrony are used such as flashbacks and flashwords known as analepsis and prolepsis respectively).

In addition to these temporal issues, the role that causality plays in describing narrative change has been relatively underexplored. As the writer and essayist George Saunders has argued, causality is the "wind in the kite" of narrative (Saunders, 2022). As Graesser et al. (2002) have demonstrated, readers are much more moved by "why" questions than "what" questions when it comes to narrative comprehension and recall. Future work will want to explore more fully both different constructs of "change" as well as draw on methodologies such as Markov models, time series analysis and systems dynamics to develop increasingly sophisticated models of change over narrative time.

Finally, most prior work is guided by a single spatial metaphor for narrative time, that of the arc. Future work will want to explore other possible structures or forms (Levine, 2015) that might capture the temporal patterns of narrative. The translation of time into spatial form represents an exciting and novel space of research for computational narrative understanding.

5 Narrative Schemas

Narratives are forms of information compression (Berns, 2022). They select certain experiential data and structure this data into prescribed grammatical slots (as described in Section 1). This basic insight serves as the foundation of the theory of narrative "scripts" (Schank and Abelson, 1977; Chambers and Jurafsky, 2008b), where narrative is understood as a probabilistic sequence of actions (i.e. given the event of being in a restaurant certain subsequent actions are more or less likely). Such compression is what allows stories to be both memorable as well as easily shareable (i.e. tellable (Baroni, 2011)).

The discussion of narrative form or shape in the prior section is one such example of the *schematic* nature of narrative, i.e. that narratives have structure and this structure is essential to their meaning. But schemas can also be represented as a variety of conceptual metaphors (that often have spatial associations). For example, in the field of clinical psychology researchers refer to two self-narrative schemas, called narratives of redemption (when bad things turn good) and narratives of contamination (when good things turn bad) (McAdams et al., 2001). Patients who structure life experience into the former schema are far more likely to be associated with positive mental health outcomes than those who engage in telling their life stories according to the latter.

The first extensive (and later controversial) study of narrative schemas emerged in the field of folklore studies (Dundes, 1962). Faced with large collections of documents with a high degree of repetitiveness, folklorists began developing systems for classifying stories according to different typologies. The most famous undertakings were Stith Thompson's Motif-Index of Folk-literature (Thompson, 1989), the Aarne-Thompson-Uther (ATU) Tale Index, and Vladimir Propp's emphasis on character "function" (Propp, 2010). Fundamental to this research was the insight that certain larger narrative patterns are maintained while local units can be changed. As Propp (2010) highlighted, whether it is an eagle or a horse or a ring that is the gift that carries away its recipient, the point of each of these stories is the event of being transported, or even more generally, the danger or affordance of gift giving.

While it is beyond the scope of this paper to rehearse debates around narrative classification (for a review see Dundes, 1962; Broadwell et al., 2018), there remains a fundamental value in developing narrative taxonomies for different domains. Narratives are indeed reducible to schemas and those schemas serve particular social and psychological functions. And yet we currently lack agreed-upon or widely used frameworks for discussing schemas, either at the individual or socio-cultural level.

Folklorist and computational narratologist Timothy Tangherlini has begun using the idea of schemas to study conspiracy theories circulating through social media (Tangherlini et al., 2020; Chong et al., 2021; Shahsavari et al., 2020), which function much like folklore in that various narrative units (Bill Gates, 5G) can be utilized for larger functional purposes (a global cabal of elites is controlling us). Related research by Mendelsohn et al. (2023) looks at "dogwhistle" detection, which can be understood as phrases with latent, toxic meanings and that likely have a narrative element to them.

Understanding schemas requires two challenging research questions. The first we can refer to as *motif tracking*, which requires the ability to model variability and repetition at both the level of local units (agents, actions, objects) and more general schemas (when certain units are deployed to tell certain kinds of stories). While systems currently exist to identify the narrative units described in Section 1 (including agents, actions, and objects), we still need ways of aggregating these units into story "types." When is Bill Gates being used to tell a story about global elites and when is he being used to tell a story about the power of philanthropy?

More importantly, we want to model the causes as well as social effects of these different story types. Do we see certain narrative schemas deployed in response to major social events (for example what are the prevalent narrative responses to financial or political or climatic shocks?). Or can certain narrative schemas predict future behavior? Similar to the clinical psychology example mentioned above but moving into the social realm, do we see the persistent invocation of narratives of national decline associated with shifts in electoral behavior? If we assume narrative is a key predictor of human behavior, we need more reliable and sophisticated ways of classifying narratives to better understand their causes and effects.

The second key dimension in studying narrative schemas is the aspect of *social dynamics*. As folklore studies first highlighted, narrative types are aggregates of numerous local instances of storytelling behavior. Each unit (whether an oral tale or social media post) may contribute to a larger narrative schema but may itself only loosely embody this schema. Narratologists refer to these local dynamics as "small stories" (Georgakopoulou, 2007), i.e. when a larger story is told through the participation of numerous actors. The quintessential example of this behavior is the "family dinner table," where family lore is the product of multiple actors engaging in the process of narrative recounting, potentially over long spans of time. At the macro-level narratologists refer to these larger narrative schemas – the aggregate of small stories – as "collective narratives" (Bliuc and Chidley, 2022), "ontological narratives" (Somers, 1994), or "deep stories" (Hochschild, 2018).

Social media and online news (broadly understood) greatly expand the complexity of collective narrative construction and small-story dynamics. One can imagine "top-down" approaches that start with known schemas and then classify individual stories or collections of stories within these taxonomies or "bottom-up" approaches that cluster individual stories into larger schemas that emerge from the collective behavior among the data. Modeling this complex, large-scale narrative behavior represents one of the major challenges for the field but one that has the most explanatory pay-offs in terms of understanding social behavior.

6 Narrative Infrastructures

As computational narrative understanding comes into its own as a distinct field within the NLP community, now is a good time to begin coordinating more of this research effort. These initiatives can take the form of shared tasks, dataset curation, and collective efforts to develop systems for narrative classification.

Shared tasks have a long history within NLP, though to date only three have been proposed for narrative understanding. The first is the narrative cloze test (Chambers and Jurafsky, 2008a; Mostafazadeh et al., 2016; Hatzel and Biemann, 2023), where systems predict the next agent-event in an event chain. Zehe et al. (2021) have proposed a task for detecting narrative scenes, while Reiter et al. (2019) have proposed a task for detecting narrative levels (when diegetic worlds are imbedded within one another, either in the form of stories within stories or temporal anachronisms such as flashbacks). Piper and Bagga (2022b) and Hatzel and Biemann (2023) have proposed annotation frameworks for narrative detection, i.e. identifying the degree to which a stretch of discourse can be identified as containing narration.

Future work will want to refine these existing initiatives as well as develop systems for the further detection of the remaining discursive units described in Section 2 (i.e. description, dialogue, evaluation). The automated identification of narrative communication in particular will prove extremely valuable for broader social and cultural analysis.

Given the value of narrative for understanding human behavior it is somewhat surprising how few datasets are available for the study of human storytelling. Much of this is due to intersecting problems of intellectual property restrictions, large library collections with low-levels of metadata, and the dynamic and ever-changing nature of online storytelling. Underwood et al. (2020) provide a large-scale annotation of ca. 200,000 fictional narratives in English in the Hathi Trust Digital Library that has been refined and updated by Bagga and

Piper (2022) to include a comparison corpus of non-fiction prose across 1.5 million sampled pages published since 1800. Hamilton and Piper (2023) extends this work to include multilingual fiction annotation across 521 different languages. Erlin et al. (2022) provide metadata on translations of fiction into English from 120 different languages also located in the Hathi Trust.

Outside of the HathiTrust, Piper (2022b) provides derived data on a collection of 2,700 works of professionally published English prose drawn from 12 different genres including Goodreads' user ratings. Mostafazadeh et al. (2016) developed an artificial corpus of very short stories (4-5 sentences) generated by crowdsourced workers. Ouyang and McKeown (2014) curated a collection of ca. 5,000 AskReddit stories told by users in response to particular prompts (e.g. what is your scariest real-life story?).

Researchers in the field should be aware that while Project Gutenberg offers a large collection of potentially narrative texts, problems of sample selection and poor metadata can lead to downstream problems that result in erroneous claims (Piper, 2022a). For addressing cultural and historical questions, researchers are strongly encouraged to use the collections described above.

Incumbent on all of these initiatives is a greater investment in inter-disciplinary collaboration. Computational narrative understanding will benefit as an endeavor with deeper collaborations between humanists and social scientists and the NLP community. As detailed in Piper et al. (2021), narratology is a field with a long and robust theoretical tradition. Those in the NLP field working on computational systems will benefit from expert collaborations with researchers who have deep backgrounds in studying narratives. Similarly, narratologists and their research frameworks stand to benefit from exposure to computational models (Piper and Bagga, 2022b). It is time to invest more heavily in these larger cross-disciplinary collaborations, especially if we aim to address the larger socio-cultural goals outlined in this paper.

7 Conclusion

As Vizenor envisioned, narratives are things we live by. They provide meaning and hold communities together. They play a role in financial, political, and psychological decision-making. The production of imaginary narratives in particular represent a massive cultural industry, spanning book publishing, movie-making, and gaming. The field of computational narrative understanding has made impressive strides in developing systems to study the causes and effects of narrative behavior across a diverse array of languages and cultural domains. We are in the process of establishing key workshops, tasks, and datasets.

By way of conclusion, I provide a sliding scale of calls to action, located from particular to general (Table 1). It is worth noting that an essential component of the field should include attention to the limiting factors of narrative, i.e. the way narratives encode experience in very particular ways and because of their persuasive power can also mislead individuals in profound ways. Greater attention to the risks of narration should therefore remain front and center as part of the endeavor of computational narrative understanding.

	Challenge Areas
Complexity ↓	1. Data Set Creation
	2. Narrative Element Detection
	3. Multilingual Modeling
	4. Multimodal Modeling
	5. Narrative Discourse Detection
	6. Narrative Time Modeling
	7. Narrative Schemas and Taxonomies
	8. Collective Stories and Social Behavior

Table 1: List of challenge areas in increasing order of generality and complexity

Acknowledgements

This research was generously supported by the Social Sciences and Humanities Research Council of Canada (435-2022-0089).

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.

Anjali Adukia, Alex Eble, Emileigh Harrison, Hakizumwami Birali Runesha, and Teodora Szasz. 2021. What we teach about race and gender: Representation in images and text of children's books. Technical report, National Bureau of Economic Research.

Taylor Arnold, Lauren Tilton, and Annie Berke. 2019.

Visual style in two network era sitcoms. *Journal of Cultural Analytics*, 4(2).

Sunyam Bagga and Andrew Piper. 2022. Hathi 1m: Introducing a million page historical prose dataset in english from the hathi trust. *Journal of Open Humanities Data*, 8(7).

Mieke Bal and Christine Van Boheemen. 2009. *Narratology: Introduction to the theory of narrative*. University of Toronto Press.

David Bamman, Snigdha Chaturvedi, Elizabeth Clark, Madalina Fiterau, and Mohit Iyyer. 2019. Proceedings of the first workshop on narrative understanding. In *Proceedings of the First Workshop on Narrative Understanding*.

David Bamman, Ted Underwood, and Noah A Smith. 2014. A bayesian mixed effects model of literary character. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 370–379.

Raphaël Baroni. 2011. Tellability. In *The living handbook of narratology*.

Aman Berhe, Camille Guinaudeau, and Claude Barras. 2022. Survey on narrative structure: from linguistic theories to automatic extraction approaches. In *Traitement Automatique des Langues, Volume 63, Numéro 1: Varia [Varia]*, pages 63–87.

Gregory Berns. 2022. *The Self Delusion: The New Neuroscience of How We Invent—and Reinvent—Our Identities*. Hachette UK.

Helena Bilandzic and Rick Busselle. 2013. Narrative persuasion. *The SAGE handbook of persuasion: Developments in theory and practice*, pages 200–219.

Ana-Maria Bliuc and Alexander Chidley. 2022. From cooperation to conflict: The role of collective narratives in shaping group behaviour. *Social and Personality Psychology Compass*, 16(7):e12670.

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.

William F Brewer and Edward H Lichtenstein. 1982. Stories are to entertain: A structural-affect theory of stories. *Journal of pragmatics*, 6(5-6):473–486.

Peter Broadwell, David Mimno, and Timothy Tangherlini. 2018. The tell-tale hat: Surfacing the uncertainty in folklore classification. *Journal of Cultural Analytics*.

Jerome Bruner. 1991. The narrative construction of reality. *Critical inquiry*, 18(1):1–21.

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.

Italo Calvino. 1988. *Six memos for the next millennium*. Harvard University Press.

Nathanael Chambers and Dan Jurafsky. 2008a. 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. 2008b. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797.

David Chong, Erl Lee, Matthew Fan, Pavan Holur, Shadi Shahsavari, Timothy Tangherlini, and Vwani Roychowdhury. 2021. A real-time platform for contextualized conspiracy theory analysis. In *2021 International Conference on Data Mining Workshops (ICDMW)*, pages 118–127. IEEE.

Multigraph Collective. 2019. *Interacting with Print: Elements of Reading in the Era of Print Saturation*. University of Chicago Press.

Laura Coltelli. 1990. *Winged words: American Indian writers speak*. U of Nebraska Press.

Alan Dundes. 1962. From etic to emic units in the structural study of folktales. *The Journal of American Folklore*, 75(296):95–105.

Lauren Edlin and Joshua Reiss. 2023. Identifying visual depictions of animate entities in narrative comics: An annotation study. In *Proceedings of the The 5th Workshop on Narrative Understanding*, pages 82–91.

Katherine Elkins. 2022. *The Shapes of Stories: Sentiment Analysis for Narrative*. Cambridge University Press.

Matt Erlin, Andrew Piper, Douglas Knox, Stephen Pentecost, and Allie Blank. 2022. The transcomp dataset of literary translations from 120 languages and a parallel collection of english-language originals. *Journal of Open Humanities Data*, 8.

Matt Erlin, Andrew Piper, Douglas Knox, Stephen Pentecost, Michaela Drouillard, Brian Powell, and Cienna Townson. 2021. Cultural capitals: Modeling minor european literature. *Journal of Cultural Analytics*, 6(1).

Elizabeth Evans and Matthew Wilkens. 2018. Nation, ethnicity, and the geography of british fiction, 1880-1940. *Journal of Cultural Analytics*, 3(2).

Gustav Freytag. 1895. *Technique of the drama: An exposition of dramatic composition and art*. S. Griggs.

Mikaela Irene Fudolig, Thayer Alshaabi, Kathryn Cramer, Christopher M Danforth, and Peter Sheridan Dodds. 2023. A decomposition of book structure through ousiometric fluctuations in cumulative wordtime. *Humanities and Social Sciences Communications*, 10(1):1–12.

Varun Gangal, Steven Y Feng, Malihe Alikhani, Teruko Mitamura, and Eduard Hovy. 2022. Nareor: The narrative reordering problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10645–10653.

Gérard Genette. 1983. *Narrative discourse: An essay in method*, volume 3. Cornell University Press.

Gérard Genette. 1992. *The architext: An introduction*, volume 31. Univ of California 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.

Jonathan Gottschall. 2021. *The story paradox: how our love of storytelling builds societies and tears them down*. Hachette UK.

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.

Sil Hamilton and Andrew Piper. 2023. Multihathi: A complete collection of multilingual prose fiction in the hathitrust digital library. *Journal of Open Humanities Data*, 9.

Md Kamrul Hasan, Wasifur Rahman, AmirAli Bagher Zadeh, Jianyuan Zhong, Md Iftekhar Tanveer, Louis-Philippe Morency, and Mohammed (Ehsan) Hoque. 2019. UR-FUNNY: A multimodal language dataset for understanding humor. 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 2046–2056, Hong Kong, China. Association for Computational Linguistics.

Hans Ole Hatzel and Chris Biemann. 2023. Narrative cloze as a training objective: Towards modeling stories using narrative chain embeddings. In *Proceedings of the The 5th Workshop on Narrative Understanding*, pages 118–127.

David Herman. 2009. *Basic elements of narrative*. John Wiley & Sons.

Jack Hessel, Ana Marasovic, Jena D. Hwang, Lillian Lee, Jeff Da, Rowan Zellers, Robert Mankoff, and Yejin Choi. 2023. Do androids laugh at electric sheep? humor "understanding" benchmarks from the new yorker caption contest. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 688–714, Toronto, Canada. Association for Computational Linguistics.

Arlie Russell Hochschild. 2018. *Strangers in their own land: Anger and mourning on the American right*. The New Press.

Peter Hühn. 2014. *Event and eventfulness*. de Gruyter.

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.

Alípio Mário Jorge, Ricardo Campos, Adam Jatowt, and Sumit Bhatia. 2019. The 2nd international workshop on narrative extraction from text: Text2story 2019. In *European Conference on Information Retrieval*, pages 389–393. Springer.

Eve Kraicer and Andrew Piper. 2019. Social characters: the hierarchy of gender in contemporary english-language fiction. *Journal of Cultural Analytics*, 3(2).

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.

Caroline Levine. 2015. Forms. In *Forms*. Princeton University Press.

Dan P McAdams, Jeffrey Reynolds, Martha Lewis, Allison H Patten, and Phillip J Bowman. 2001. When bad things turn good and good things turn bad: Sequences of redemption and contamination in life narrative and their relation to psychosocial adaptation in midlife adults and in students. *Personality and social psychology bulletin*, 27(4):474–485.

Laura B. McGrath. 2018. *Middlemen: Making Literature in the Age of Multimedia Conglomerates*. Ph.D. thesis, Michigan State University.

Dane Malenfant Margaret Meehan and Andrew Piper. 2022. Causality mining in fiction. In *Proceedings of Text2Story-Fifth Workshop on Narrative Extraction From Texts held in conjunction with the 44th European Conference on Information Retrieval (ECIR 2022), Stavanger, Norway*, volume 3117, pages 25–34.

Julia Mendelsohn, Ronan Le Bras, Yejin Choi, and Maarten Sap. 2023. From dogwhistles to bullhorns: Unveiling coded rhetoric with language models. *arXiv preprint arXiv:2305.17174*.

Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849.

Elinor Ochs, Lisa Capps, et al. 2009. *Living narrative: Creating lives in everyday storytelling*. Harvard University Press.

Jessica Ouyang and Kathleen McKeown. 2015. Modeling reportable events as turning points in narrative. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2149–2158, Lisbon, Portugal. 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).

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.

Charuta Pethe, Allen Kim, and Steven Skiena. 2020. Chapter captor: Text segmentation in novels. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8373–8383.

Federico Pianzola. 2018. Looking at narrative as a complex system: The proteus principle. In *Narrating complexity*, pages 101–122. Springer.

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. 2022a. Biodiversity is not declining in fiction. *Journal of Cultural Analytics*, 7(3).

Andrew Piper. 2022b. The conlit dataset of contemporary literature. *Journal of Open Humanities Data*, 8.

Andrew Piper. 2023a. Modeling narrative revelation. *Under review*.

Andrew Piper. 2023b. What do characters do? the embodied agency of fictional characters. *Journal of Computational Literary Studies*, 2(1).

Andrew Piper and Sunyam Bagga. 2022a. A quantitative study of fictional things. *Proceedings of the Computational Humanities Research Conference, December 12 – 14, 2022, Antwerp, Belgium*, pages 268–279.

Andrew Piper and Sunyam Bagga. 2022b. Toward a data-driven theory of narrativity. *New Literary History*, 54(1):879–901.

Andrew Piper, Richard Jean So, and David Bamman. 2021. Narrative theory for computational narrative understanding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 298–311.

Andrew Piper and Olivier Toubia. 2023. A quantitative study of non-linearity in storytelling. *Poetics*, 98:101793.

Vladimir Propp. 2010. *Morphology of the Folktale*, volume 9. University of Texas Press.

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*, 4(3).

Paul Ricoeur. 2012. *Time and Narrative, Volume 1*. University of Chicago press.

Marie-Laure Ryan, James Ruppert, and John W Bernet. 2004. *Narrative across media: The languages of storytelling*. U of Nebraska Press.

Christian Salmon. 2017. *Storytelling: Bewitching the modern mind*. Verso books.

Brenda Santana, Ricardo Campos, Evelin Amorim, Alípio Jorge, Purificação Silvano, and Sérgio Nunes. 2023. A survey on narrative extraction from textual data. *Artificial Intelligence Review*, pages 1–43.

Maarten Sap, Anna Jafarpour, Yejin Choi, Noah A Smith, James W Pennebaker, and Eric Horvitz. 2022. Quantifying the narrative flow of imagined versus autobiographical stories. *Proceedings of the National Academy of Sciences*, 119(45):e2211715119.

George Saunders. 2022. *A Swim in a Pond in the Rain*. Bloomsbury Publishing.

Roger C. Schank and Robert P. Abelson. 1977. *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. L. Erlbaum, Hillsdale, NJ.

Benjamin M Schmidt. 2015. Plot arceology: A vector-space model of narrative structure. In *2015 IEEE International Conference on Big Data (Big Data)*, pages 1667–1672. IEEE.

Shadi Shahsavari, Pavan Holur, Tianyi Wang, Timothy R Tangherlini, and Vwani Roychowdhury. 2020. Conspiracy in the time of corona: automatic detection of emerging covid-19 conspiracy theories in social media and the news. *Journal of computational social science*, 3(2):279–317.

Robert J Shiller. 2020. *Narrative economics: How stories go viral and drive major economic events*. Princeton University Press.

Margaret R Somers. 1994. The narrative constitution of identity: A relational and network approach. *Theory and society*, pages 605–649.

Dominik Stammbach, Maria Antoniak, and Elliott Ash. 2022. Heroes, villains, and victims, and gpt-3– automated extraction of character roles without training data. *arXiv preprint arXiv:2205.07557*.

Meir Sternberg. 1990. Telling in time (i): Chronology and narrative theory. *Poetics Today*, 11(4):901–948.

Meir Sternberg. 1992. Telling in time (ii): Chronology, teleology, narrativity. *Poetics today*, 13(3):463–541.

Timothy R Tangherlini, Shadi Shahsavari, Behnam Shahbazi, Ehsan Ebrahimzadeh, and Vwani Roychowdhury. 2020. An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, pizzagate and storytelling on the web. *PloS one*, 15(6):e0233879.

Stith Thompson. 1989. *Motif-index of Folk-literature*. Indiana University Press.

William HW Thompson, Zachary Wojtowicz, and Simon DeDeo. 2018. Lévy flights of the collective imagination. *arXiv preprint arXiv:1812.04013*.

Olivier Toubia, Jonah Berger, and Jehoshua Eliashberg. 2021. How quantifying the shape of stories predicts their success. *Proceedings of the National Academy of Sciences*, 118(26):e2011695118.

Ted Underwood. 2018. Why literary time is measured in minutes. *ELH*, 85(2).

Ted Underwood, Patrick Kimutis, and Jessica Witte. 2020. Noveltm datasets for english-language fiction, 1700-2009. *Journal of Cultural Analytics*, 5(2).

Michael Vauth, Hans Ole Hatzel, Evelyn Gius, and Chris Biemann. 2021. Automated event annotation in literary texts. In *CHR*, pages 333–345.

Piek Vossen, Tommaso Caselli, and Roxane Segers. 2021. A narratology-based framework for storyline extraction. *Computational Analysis of Storylines: Making Sense of Events*, 125.

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.

Gregory Yauney, Ted Underwood, and David Mimno. 2019. Computational prediction of elapsed narrative time. *Workshop on Narrative Understanding*.

Albin Zehe, Leonard Konle, Lea Katharina Dümpelmann, Evelyn Gius, Andreas Hotho, Fotis Jannidis, Lucas Kaufmann, Markus Krug, Frank Puppe, Nils Reiter, et al. 2021. Detecting scenes in fiction: A new segmentation task. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3167–3177.