

Creative Commons Attribution 4.0 International (CC BY 4.0)

<https://creativecommons.org/licenses/by/4.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

**Please provide feedback**

Please support the ScholarWorks@UMBC repository by emailing [scholarworks-group@umbc.edu](mailto:scholarworks-group@umbc.edu) and telling us what having access to this work means to you and why it's important to you. Thank you.

Received 18 July 2022, accepted 1 September 2022, date of publication 8 September 2022, date of current version 30 September 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3205314

## TOPICAL REVIEW

# Computational Understanding of Narratives: A Survey

PRIYANKA RANADE<sup>ID</sup>, SANORITA DEY, ANUPAM JOSHI<sup>ID</sup>, (Fellow, IEEE), AND TIM FININ<sup>ID</sup>

Department of Computer Science & Electrical Engineering, University of Maryland, Baltimore County, Baltimore, MD 21250, USA

Corresponding author: Priyanka Ranade (priyankaranade@umbc.edu)

This work was supported in part by NSA and in part by the National Science Foundation under Grant 2114892.

**ABSTRACT** Storytelling, and the delivery of societal narratives, enable human beings to communicate, connect, and understand one another and the world around them. Narratives can be defined as spoken, visual, or written accounts of interconnected events and actors, generally evolving through some notion of time. Today, information is typically conveyed over online communication mediums, such as social media and blogging websites. Consequently, the act of narrative delivery itself has shifted from simply imparting information through self-contained structures such as books, to more fragmented structures, such as social media websites, where evolving story events are constructed over multiple online sources. Ubiquitous online conversation can manifest into sophisticated narratives that have the potential to influence wide-spread user interpretations of cultural sentiments, attitudes, values, as well as geopolitical events and facts. As a result, narratives are actively being used as strategic tools for shaping local events, promoting collective opinions, and asserting ideologies and propaganda, making them sources of interest for identifying themes, intentions, and goals across multiple communities and potential adversaries. Identifying fragmented narratives, extracting thematic and temporal components that constitute evolving narratives, and locating signs of active rhetoric framing tactics, are difficult to detect and analyze without large-scale automation. This problem can be addressed through the use of natural language understanding technologies. Our goal is to document and discuss methods to efficiently construct, extract, and detect evolving online narratives. The novel contribution of this paper is the formal collation and documentation of such technologies and research areas, as well as extensive discussion on open research challenges and goals in the definition, identification, construction, generation, and representation of online narratives. To our knowledge, there is currently no existing formal documentation that organizes and provides extended discussion on narrative understanding research areas and open challenges.

**INDEX TERMS** Narratives, storytelling, narrative extraction, story evolution, narrative generation, computational semantic analysis, natural language understanding, misinformation, propaganda, journalism.

## I. INTRODUCTION

The existence of *narratives* (accounts of series of interconnected events that fulfill a *story*) and the act of *story-telling*, have been critical elements in the creation of universal human experiences across three distinct periods of communication in history, in addition to predated written account [90]. Narratives have shaped the ways in which humans have created languages to communicate and connect with one another

(the era of oral tradition), record and make sense of their surroundings (the age of literacy), to finally formalizing the ability to spread and convey information at a large scale through electronic media (the age of information) [40]. In particular, electronic media has revolutionized the methods that allow users to build online communities and easily share wide-spread information.

As a result, there is a large amount of data available on the web today, from images and videos, to unstructured and semi-structured text data, giving us the ability to share narratives at a large scale, through multiple communication

The associate editor coordinating the review of this manuscript and approving it for publication was Fu Lee Wang<sup>ID</sup>.

modes [1], [33], [137]. The high availability of data, and the process of online narrative analysis, can be beneficial for a variety of purposes, such as ensuring safe online environments for sharing consistent and factual information, and providing critical intelligence to a multitude of geopolitical defense agencies. A large portion of the narratives shared online today have malicious intentions and are rooted in propaganda-spreading campaigns. Digital media outlets and advancing computational methods have heightened the sharing of biased information, misinformation, and disinformation rapidly and widely. Many of these campaigns aim to craft targeted narratives that can be spread to large audiences, with the goal of inducing high engagement such as opinionated and passionate conversation, and increased interactions such as post reactions and shares. News media outlets can further circulate deceptive messages by indiscriminately re-posting and reporting malicious campaign-led information [5]. Due to these critical, omnipresent implications, identifying both existing and evolving narratives continues to be of concern and a growing research problem. We are interested in computational methods that can automatically gather and chain information crucial to evolving narratives (construction), as well as detect and extract existing narratives.

Developments in Natural Language Understanding (NLU) continue to refine the methods that allow computational systems to capture and understand human experiences, such as the narratives we share. Much of this technology contributes to the growth of integrated systems that aim to understand, extract, and generate a variety of *self-contained* or *fragmented* narrative structures. Self-contained narratives, such as books or films, do not rely on external sources to convey the set of events that form a *complete* story, whereas fragmented narrative structures such as the collection of *chained* social media posts, blogs, or articles, are considered stories that are *constructed* using events spread across disparate sources [112].

Though there are ongoing challenges and open research questions associated with the construction, understanding, and generation of both self-contained and fragmented narratives, there is a considerably larger focus on developing methods to address fragmented narrative-based problems. This imbalance can be attributed to the nature of the ways in which human beings consume narratives in the current Information Age. While stagnant, self-contained narratives are still popularly used for both entertainment and educational purposes, users often lean towards sources such as social media sites like Twitter<sup>1</sup> and Reddit<sup>2</sup> to learn about information associated with current events that are pervasive in their surrounding societies. Fragmented narratives naturally emerge through these large information sources, and the events of the narrative (especially when considering critical,

popular news topics) are typically reported and *instantiated* by millions of users [19], [108], [149].

The sequence of events that construct fragmented narratives can become very detached, due to the difficulties in accurately and efficiently locating evolving themes, concepts, and gaps in ongoing narrative chains [45], [111], [113], [159]. An example of the construction of a fragmented narrative about the *January 2022 Texas Synagogue Hostage* from multiple online sources is shown in Figure 1.

In this scenario, breaking events and details about the hostage are extracted from multiple disparate online sources, and *constructed* chronologically based on their context. The ordering of major events results in a logical, end-to-end narrative about the hostage. The methods required to first gather pieces of information and later construct them in logical storylines, span a multitude of research areas in Information Retrieval (IR) and Natural Language Understanding (NLU). The first step is data acquisition, more specifically, locating and extracting meaningful and *relevant* news events, with the extended challenges of closing potential gaps in a story chain of events. Temporal IR methods such as measuring relevance of former content, article summarization, dynamic topic modeling, and information filtering, are examples of methods used in the news event extraction phase. The chronological ordering of events into plot sequences is also known as *narrative construction*. Examples of methods used to address narrative construction research problems include shift detection, and causal relation extraction and arrangement.

Narrative construction differs from *automated journalism*, which is a form of computer-generated news generation, where news stories are automatically produced by machines rather than human journalists [9]. The goal of narrative construction is to enhance traditional information retrieval methods through typical search engines. Though search engines can retrieve aggregated event information, they lack the ability to sequence relevant events together to form narratives about different topics. Unlike automated journalism, narrative construction does not generate original content through automated methods, but rather chains insights from existing articles written by human reporters. The goal is to provide *ordered* sequences of events to end users, enhancing the traditional aggregated output that typical search engines provide.

The novel contribution of our work is the collation, organization, and documentation of multiple research areas, shown in Figure 2. We define structural components of universal narratives, as well as expand on specific research methods and projects to extract, construct, generate, and represent narratives. It should be noted that there are existing focused workshops that have created communities for sharing research in the areas of narrative generation, understanding, and construction [14], [15], [39], [66], [67]. However, there is a lack of a single resource that organizes and describes fundamental concepts related to defining common structural components of narratives, as well as, foundational previous work, current research areas, and future directions. Our research is inspired and guided by the following research questions:

<sup>1</sup><https://www.twitter.com>

<sup>2</sup><https://www.reddit.com/>

*‘RQ1: How do we differentiate between online and traditional narratives?’*

Before elaborating on the specific projects that fall under each research area (Section IV), we first describe and define different types of narratives, more specifically, distinguishing between self-contained and fragmented narratives (Section III-A). Today, narratives can be classified into several categories, but are typically generalized in three overall formats: visual stories (movies, videos), audible stories (audiobooks, podcasts, news), and written stories (books, articles, blogs, social media posts).

*‘RQ2: What primary components do narrative structures contain?’*

Though the *methods* each of the three narrative formats use to actually communicate stories differ tremendously from one another, the comprehensive concept of the underlying narrative itself is consistent across all three. For each of these forms of narratives, the story is developed through the execution of the relationships and information about *and* between people, surrounding places and objects, and the events that encompass them over time. Section III-B outlines major components of a narrative (entities, semantic relationships, and plots).

*‘RQ3: What are existing computational methods that have the ability to represent and validate narrative structures?’*

Once the narrative types and components are introduced, computational methods to represent and validate narrative structures can be thoroughly explored. Section IV thoroughly outlines existing work and open challenges in the research areas shown in Figure 2.

More information on the motivation for each of the research questions, establishment of primary research areas, as well as the methods we used to collect and organize the research areas and relevant projects is described in Section II. A summary of key research projects for each of the research areas pictured in Figure 2 is documented in Table 3.

More specifically, this paper has the following novel contributions:

- Definition of the structural components of narrative structures, thoroughly describing each narrative component and the existing relationships between them.
- Collation of existing IR and NLU-based approaches into an extensive and clear research area hierarchy.
- Documentation of key research projects in narrative extraction, generation, story evolution, evaluation, and existing resources and tools.
- Identification of the focus, contributions, and existing challenges present in each current area.
- Discussion of growing research trends and areas for future research.

The rest of this paper is organized as follows. We begin with Section II, by describing the research questions that motivated our research and providing context around each of the research areas (Figure 2) that address the research questions. We also include the methods that we used to shortlist

the primary research areas and specific research projects that fall under each area. Once the narrative types and components are introduced, Section IV thoroughly outlines existing work and open challenges in the research areas shown in Figure 2. We conclude in Section V and give examples of future research directions surrounding narrative construction and understanding tasks.

## II. RESEARCH AREA MOTIVATIONS AND SURVEY COMPILATION METHODS

Given the critical role narratives play in human interaction and assimilation, we were curious about the patterns in which narratives are *developed* online, as well as motivated by the potential for computational systems to process and understand evolving narratives. The proliferation of digital technologies has increased the number of mediums for human beings to share information widely and efficiently. This shift in communication presents increased opportunities to observe human level interaction at a magnified scale, but also brings a number of challenges for managing and extracting meaningful insights from the abundance of available online data. This motivates us to explore computational methods to improve the detection, construction, and generation of narratives via automated methods.

### A. RESEARCH GOALS

We use the research questions described in Section I to guide the literature review.

Our first goal is to provide documentation of narratives through a high level viewpoint by differentiating between types of narratives, and more specifically, distinguishing between traditional, self-contained narratives, and more novel emerging formats such as fragmented narratives. Along with documenting the diversity of narratives present today, we also aim to understand components that are universal to the majority of narratives in order to make the research problem more attainable for computational representation and analysis. This universal view of narrative types and their components guided our initial literature survey on the origin of narratives, existing narrative types, and linguistic components that have the ability to be processed using methods in computational semantic analysis.

We further explored methods rooted in Natural Language Understanding (NLU) to address various challenges in narrative construction and understanding. For example, we were curious about potential methods used to automatically differentiate between narrative types and extract the narrative components we resolved from RQ2. As discussed previously in Section I, we are mainly motivated by narratives that develop over fragmented online sources. As such, we conducted an additional extensive literature survey on methods that involve creating story-chains from multiple sources, and constructing events in chronological timelines.

Though narrative extraction and construction are the two most obvious areas to address the fragmented narrative research problem, we additionally conducted a survey

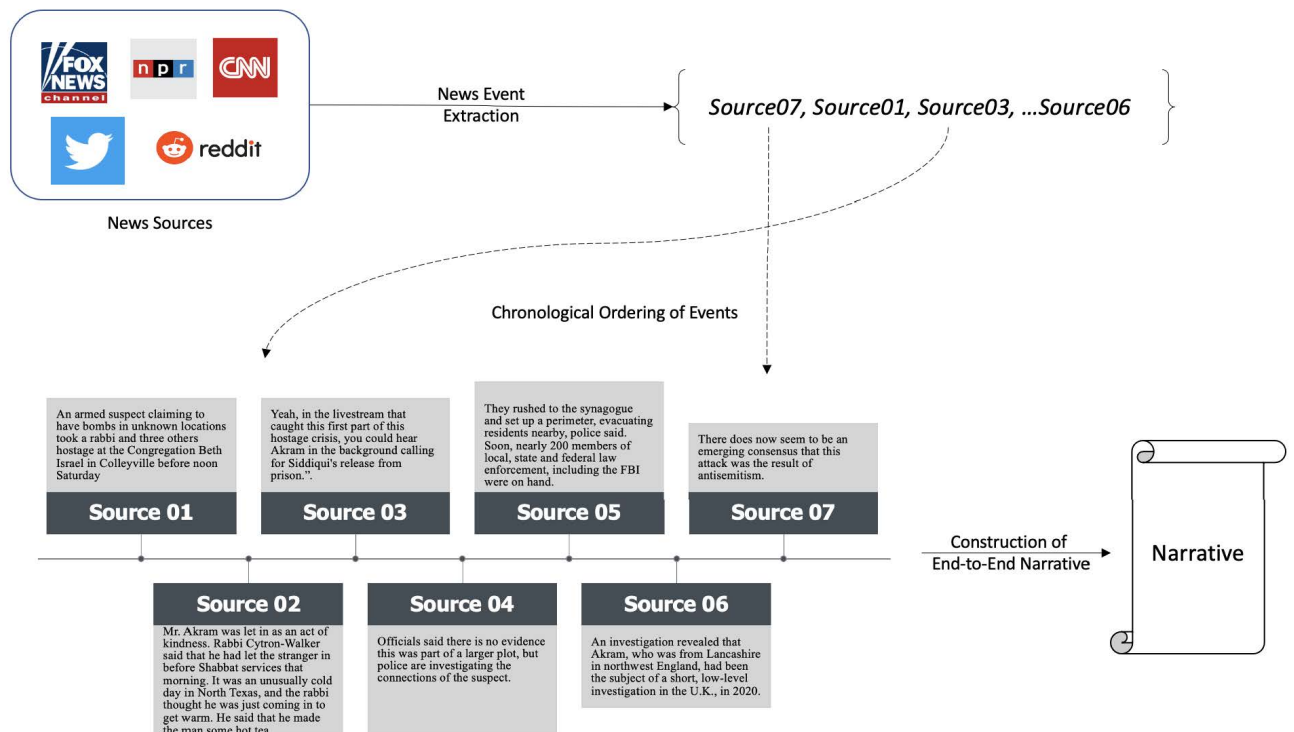


FIGURE 1. Construction of fragmented narrative about the 2022 Texas synagogue hostage.

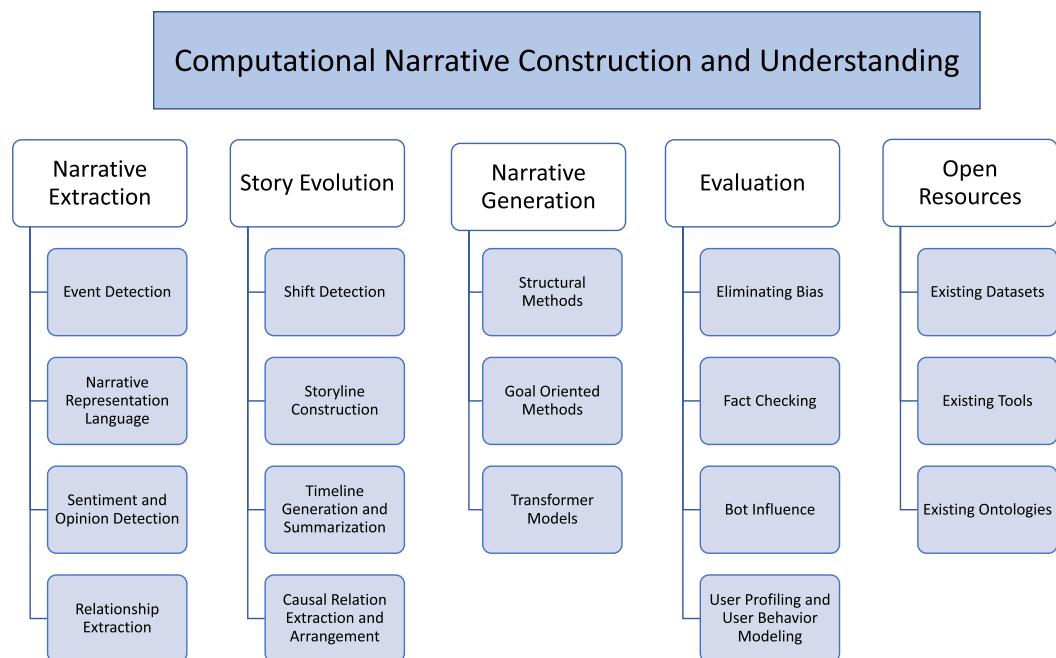


FIGURE 2. Roadmap of research opportunities in narrative construction and understanding.

on narrative generation methods. One of our goals is to explore whether generative methods can reveal inherent rhetorical tactics such as *framing* within input sequences.

More specifically, we were curious to understand whether the ordering of processed words, can influence the interpretation of future outcomes in an ongoing narrative.



Lastly, the novelty of this paper is providing a communal base of understanding and documentation for both narrative types, as well as, the existing computational methods used to understand narratives. In accordance to this goal, it is important to also document available open resources to continuously and communally extend. Some of these resources can include potential ontologies, data models, evaluation baselines, and training data. Our methods for collecting the sources included as part of this survey are provided in the next section.

## B. LITERATURE SURVEY METHODS

There is no current resource to our knowledge that has compiled narrative analysis research areas and specific papers into a single document. Therefore, we provide this section to validate the comprehensiveness of our survey research, describe the inclusion and exclusion criteria, and elucidate the query methodology used to select the concepts, topics, and research work for this survey.

In order to devise the research area hierarchy presented in Figure 2, we first located a series of existing workshops and journals with a focus in narrative understanding and storytelling [14], [15], [39], [66], [67]. These workshops were developed by researchers who have made significant contributions to the narrative understanding field. We utilized their subject matter expertise to organize relevant, high level topics of interest to create a shortlist of the most critical research areas to thoroughly document in this paper. Examples of topics include “narrative extraction” and “storyline evolution”. The final list of topics are displayed as the white boxes in Figure 2. For each workshop, we observed the changes in topics for each existing year, and also included additional topics as technologies and priorities changed. We also excluded topics that did not persist past a single workshop and journal year, or across multiple workshops and journals. We used these topics as search terms to query relevant workshops and papers.

Once we surveyed persistent and relevant topics in the field, we formulated a hierarchy of research areas (Figure 2) and used them to guide our search queries for sub-topics. We shortlisted papers on the basis of a combination of pioneering and novel ideas, large number of citations, and proven reproducibility of methods. Our goal is for other researchers in the narrative understanding community to refer to the hierarchy of research areas to help guide future research questions and directions.

In the next section, we address research questions, RQ1 and RQ2, by documenting types of narratives and components of universal narratives.

## III. DEFINING NARRATIVE STRUCTURES

### A. TYPES OF NARRATIVES

This section addresses RQ1 (defined in Section I). Narratives can undertake many formats and are expressed through various types of media. The most colloquial and

traditional format of a narrative is a *literary narrative*, where the story is conveyed in the format of a novel (physical or electronic-book), which contains written word [43]. However, literary novels are not limited to written word, and can also contain illustrations, or even at times, strictly illustrations (no written word), better known as a form of *visual narrative* [43], [114]. Other media formats that convey visual narratives include movies and videos [51]. There have been several research projects done on understanding visual narratives [6], [59], [79].

This paper, however, focuses on the survey of the computational understanding and extraction of *written narratives*. Other examples of written narratives, though not as communally understood as such in comparison to literary narratives, include physical or electronic newspapers [9], social media posts [112], or collections of online blogs [11]. As stated previously in Section I, there is a larger focus on computational understanding and extraction of *fragmented narratives*, which are typically stories divided across collections of newspapers, blogs, and social media posts, and are not statically defined within *self-contained* stories. A fragmented narrative is a construction of stories where the narration itself is inherently dynamic, constantly shifting the story events as well as the intended outcome [159].

When considering online blogs and social media websites, narratives can also be used to convey ideologies and opinions [78]. Consequently, users of these media outlets can obtain influenced views (intentionally or unintentionally) of their societies, through fragmented accounts of cultural sentiments, attitudes, and values [25]. In addition, users can also gain insight into views of additional society models through the narratives they may expose themselves to [25]. Narratives can therefore be used as *strategic tools*, to shape local events and promote collective opinions [25], [50]. Narratives can also be further divided into *domain-narratives*, stories that form within specific fields of study, such as finance [56], the gaming industry [21], and clinical sciences [124]. More discussion on domain-narratives is available in Section IV-A. Similar to the concept of domain-specific narratives, there are also thematic classifications on types of narratives, also known as *story tropes* [125].

Often, narratives can be part of multiple domains simultaneously, making it a difficult task to identify components of the narrative accurately and efficiently. Identifying the narrative, as well as extracting thematic and temporal segments of the narrative are challenging research problems. This paper aims to identify open research areas in identifying, constructing, and extracting narratives found across multiple sources. In summary, segments of a narrative can be conveyed across several sources and periods of time, making it a difficult task to identify the emergence of a narrative, identify gaps in a storyline, or piece together accurate events of a narrative.

Before reviewing the current research areas as shown in Figure 2, it is important to understand the fundamental components that generally form most narrative structures. These

components are defined and described in the next Section (Section III-B).

## B. COMPONENTS OF A NARRATIVE

Despite the diversity of narrative types that exist today, they can be generalized based on a common set of structural components described in this section. Through an extensive survey of various types of existing narratives and respective structures, we were able to resolve three common elements, of *entities*, *semantic relationships*, and *plots*, addressing RQ2 as defined in Section I.

### 1) ENTITIES

*Entities* are real-world subjects or objects such as person, time, or location, and are often denoted with a name (Named Entities) [103]. Named Entities can be further abstracted into entity instances [103]. For example, *Washington DC* is an instance of the entity *city*. Identifying entities in narratives is a crucial first step to deriving semantic meaning from a given story or a combination of stories, since story points (thematic and temporal plots) develop around entities such as persons, events, and times [109].

Named Entity recognition (NER) models are commonly used to automatically identify entities in both structured and unstructured text documents [77]. Text-based narratives are largely considered examples of unstructured textual resources since they lack *predefined formats* and are often either collections of native text or a conglomeration of information derived from a variety of formats and resources. NER models use both linguistic grammar-matching techniques as well as statistical-based models, to classify named entities in unstructured text. In addition, they typically require supervised or semi-supervised techniques for learning [38], [85], [156]. The training set for an NER model typically consists of mappings of words and their part-of-speech (POS) tags, syntactic chunk tags, and named entity tags, with the objective of teaching the model to learn mappings for unseen words given a free-text input sample [55], [121], [126], [145].

NER models trained on general entity recognition datasets are rarely domain-transferable, creating the need for domain-specific training data creation and learning [64]. Narratives exist across a variety of subjects, and as a result, have the potential to incorporate unique, domain-specific entity labels. Domain-specific narratives have started to become explored, and are explained in Section III-A.

Despite the potential of the occurrence narrative-specific entities, there are three common entities that provide a general structure to a majority of narrative extraction, construction, and understanding tasks: *actors*, *events*, and *times*. These entities can be further broken down into labels specific to a particular narrative [147]. Generally, narratives include a set of persons involved in the execution of a story (actors). In addition, actors participate in connected occurrences of contextually relevant incidents (events) that provide flow to a story over chronological periods (times). *Actors* and *times* are defined similarly across different narratives. There

is little flexibility in defining entities that represent people, or moments in time since they are more easily perceived and understood unanimously (less room for abstract meanings of components that constitute a person or a time). On the other hand, events are more difficult to formalize because researchers define events differently, depending on the narrative context. Current datasets that reflect differences in events are described in Section IV-E, and work done in the area of event detection is available in Section IV-A.

### 2) SEMANTIC RELATIONSHIPS

Entity recognition alone cannot reflect chronological succession of events, relations between events and actors, or transformations actors portray over time. For example, referring to the following sentence,

*Washington, D.C. is the capital of the United States of America.*

The meanings of the terms “Washington”, “Capital”, and “America” are limited to their individual (and potentially multiple) definitions. Without defining *relationships* between the entities, we cannot obtain the overall meaning of the entire sentence, but only the individual definitions of each word. This is especially important for tasks that aim to derive understanding *across* narratives, because temporal themes and plot points of stories develop based on a compound of persons, events, and times (entities) in addition to, the relationships, interactions, or associations that exist between them.

Extracting relationships between entities to obtain meaning from text is known as *semantic analysis*, where each link between entities is known as a *semantic relationship* [127]. There are several existing relationship extraction approaches that are able to automatically identify semantic relationships that exist between entities in a given input text [7], [82].

There are a number of different semantic relationships that have been identified by linguists, psychologists, and computer scientists studying natural language [53], [70], [136], [141]. In addition, there exist a number of structures such as hierarchical clustering diagrams and taxonomies, which organize and represent the types of relationships and sub-relationships that can exist between entities [17], [141]. Chaffin *et al.* provide an analysis of the similarities and differences in semantic relationships. The authors create a tree-like structure to decompose thirty one types of semantic relationships [17]. The number and type of semantic relationships to extract varies by the domain and annotation tasking goals for different corpora. For example, in a domain like *biomedical-engineering*, in-domain semantic relationships such as *protein-organism-location*, were identified to extract the most relevant relationships given domain-specific tasking [89]. Similarly, domain-specific relationship extraction has been performed on domains such as cybersecurity [116], [117].

Though there are variations, semantic relationships can generally be grouped under the following relationships listed in Table 1. The final relationships were identified through an

TABLE 1. Semantic relationship examples.

Semantic Relationship Type	Definition	Predicate Examples
Inclusion	One entity type inherits, contains or contains components of other entity types. Inclusion relationships include: <i>class</i> , <i>meronymic</i> , <i>topological</i> .	is-a-type-of is-part-of is-in
Attribution	Asserting or assigning ascription between entities.	is-the-author-of is-related-to is-employed-by
Temporal / Sequential	Relationship between a temporal duration-based entity and other related entities such as events or actors.	is-before is-after is-late is-early
Causal	Relationship between two entities, where one is identified as the cause and the second identified as the effect.	in-order-to because-of as-a-consequence-of
Comparative	Relation between two entities, where one entity is identified as a comparator and the other is identified as the comparable.	is-smaller is-larger is-more is-less
Possession	Asserting ownership or control between entities.	has-object

extensive review in both linguistics-based literature, including computational linguistics [17], [72], [141]. As discussed previously, the number and diversity of semantic relationship types presents the challenge of representing and capturing various relationships across multiple domains. One method, is capturing these relationships using taxonomies, which model part-whole relations between entities [150]. However it is clear that taxonomies only represent *class inclusion-based* relationships, even though knowledge about a particular domain is understood as a combination of various semantic relationships. Class inclusion-based relationships can actually branch into other relationships such as *causal* and *comparative*, creating a complex network of relationships that represent knowledge about a particular domain or sets of domains. This complex network of knowledge is known as an *ontology*, and is used to represent the connection between different semantic relationships [49]. There have been several ontologies defined for narrative-based problems and are later described in Section IV-E.

### 3) PLOTS

Using semantic relationships between entities to model the sequence of interconnected *actors* and developed *events*, over several periods of *time*, can lead to the discovery of *plots* [32]. In particular, plots represent the contextual *development* of events and actors over time within either a single narrative or a hierarchy of sub-narratives [16], [41]. Events and actors are connected to each other through the identification and reasoning of the various semantic relationships described in the previous section.

Plots themselves can be more explicitly defined using the following five elements defined in Table 2 [32], [47]. Considering the elements, it is evident that plots include the most significant events in a given story, which include changes in actors and their actions, over time. Similar to the representation of semantic relationships mentioned in the previous

TABLE 2. Defining elements of a plot.

Plot Element	Definition
Exposition	Introduction of characters, setting, and main story objective.
Rising Action	Series of events that build on main story objective.
Climax	Occurrence of the major event of the story.
Falling Action	Series of events impacted by climax.
Resolution	Events that conclude the story.

section, narrative plot elements can also be represented using ontologies. More information on ontologies for narrative plot elements is described in Section IV-E.

The rest of this paper documents research areas that focus on the understanding, extraction, and generation of narrative elements, as well as respective evaluations and open resources. For reference to the specific areas covered, refer to Figure 2.

## IV. RESEARCH AREAS

This section provides an overview of RQ3, which guided our survey on existing computational methods for narrative understanding. The research areas we surveyed are available in Figure 2, and a summary of the various research projects that fall under research area and their respective open research challenges and limitations is available in Table 3.

There is a large range of research that focuses on locating and extracting existing and evolving narratives in both self contained and fragmented formats, and is described in Section IV-A. Narratives that evolve over multiple sources add increased complexity in identifying and extracting story plots. Time is an important factor when studying the progress of narratives in both self-contained and fragmented formats. We describe work done in the tracking and extracting the evolution of a narratives in Section IV-B. A large research area also involves generating narratives automatically. We describe work using both traditional and state-of-the-art methods used for automatic narrative generation in Section IV-C. Extracting evolving narratives and generating narratives also require extensive evaluation methods to validate the primary structural components. We describe existing methods and tools used for evaluating extracted and generated narratives in Section IV-D. Lastly, we describe open tools and datasets developed for both evaluation, as well as detection, extraction, and generation of narratives in Section IV-E.

### A. NARRATIVE EXTRACTION

Locating and extracting existing narratives in self-contained or fragmented sources is a crucial research component in narrative understanding research. Narrative extraction as a research area can be further divided into three main areas of event detection, narrative representation, and sentiment and opinion detection.

As described in Section III-B, *events* are types of entities that describe main incidents that take place between actors, as a result of the actions of actors, as well as circumstances



that arise in relation to the overall setting, and are typically the building components of plot structure components shown in Table 2. Similar to the general form of detecting entities (Named Entity Recognition), there are several examples of work done in the area of specifically detecting events in large datasets of unstructured text, with the goal of identifying interconnected incidents that ultimately make up a narrative plot.

Events can further be categorized into several subtypes [138], [151]. Xiang *et al.* provide a survey of task definitions, data sources and performance evaluations for event extraction from text and create a taxonomy of approaches for different event subtasks [151]. A popular approach to identifying hierarchical events is using frame semantic parsers to detect events in unstructured text [27]. Semantic frame parsing is a language understanding task that defines a *frame* as a sentence-level event or scenario, and defines *frame elements*, as the corresponding elements and roles that can be associated to frames. Spiliopoulou *et al.* [138] use frame semantic parsers to detect event nuggets (semantically meaningful units of text that denote some action), as well as classification of types and subtypes of different types of events by creating mappings from an event task called ACE<sup>3</sup> to FrameNet taxonomy.<sup>4</sup>

There are several schemas similar to ACE, that represent hierarchical events. Many of these schemas are described in greater detail in Section IV-E. Similar to the previous paper described, Rehm *et al.* [122] also identify a class of events called Movement Action Events (MAEs), training on ACE Multilingual 2005 data. More information on the ACE dataset, and its uses for event extraction tasks, is described in Section IV-E. MAEs are defined as are entities in a sentence that refer to events involving participants and locations. The authors identify MAEs in a domain-specific task of generating travelogues (sequences of travel events that create a trip). The main approach was processing multiple interconnected instances in order to generate one instance of a travelogue, which can be applicable to other narrative extraction problems with the goal of generating storylines or plots. Event detection is also popularly done using deep learning approaches [20].

Though *events* are crucial components that build a plot or storyline, general entity recognition and extraction is also an important task in understanding surrounding elements that are related to major plot events. As described in Section III-B, there are a variety of entity types that can be identified and it is a wide research problem in selecting relevant entity types for varying narrative structures. Oza *et al.* [109] model and study different types of semantic links (relevant co-occurrence graph, unweighted link patterns, and relevance-weighted bibliographic coupling) present between entities and identify most useful relationships for a given entity retrieval task.

Similar to tasks that detect relevant structural components (semantic links and entity types) within particular narratives, another large research area involves representing primary components of a narrative through a variety of schemas or *narrative representation languages*. For example, Hussain *et al.* develop a tree-like structure to visualize narratives and prominent keywords they contain [60]. In this work, the authors focus efforts on visualizing extracted narratives from a corpus of large, unstructured text. Similar to narrative element visualization, representing narrative components is also an open research area. Yan *et al.* create functional story schemas that capture latent, structural features in *Reddit* conversation threads [154]. The authors create structured representations that surround specific narrative elements, both event and character-centric. Similarly, Labutut *et al.* survey methods for modeling relationships between narrative actors, focusing on techniques that extract character networks [75].

Relevant narrative elements can differ depending on the domain. There is several research that focuses on extracting domain-specific narratives [46]. Additional examples of domain-specific narratives are described in Section III-A.

An important aspect of identifying a narrative, especially for fragmented narratives that develop over multiple sources, is the method of tracking the evolution of a narrative, or identifying a *storyline* that binds narrative pieces together. *Story Evolution* can be its own research area, and is further documented in the next section (Section IV-B).

## B. STORY EVOLUTION

Identifying the evolutionary trends that temporally connect narrative components is known as story evolution detection. Story evolution detection can be further divided into four main research areas as shown in Figure 2. *Shift Detection*, listed as the first of these sub-areas, is a fundamental procedure that identifies transitions between plot elements. For more information on plot elements, refer to Table 2. Shift Detection is particularly important for identifying sequences in fragmented narrative structures and is a common strategy for detecting story components across a variety of sources. Hussain *et al.* [61] study shifts in narratives across social media blogs. In particular, the authors crawled news articles from blog sites related to migrant issues in Europe and analyzed trends in topics across a one year period. They further use the topic shift analysis to study sentiments of particular narratives for a given topic. Similarly, Marcoux *et al.* [95] study the chronological themes of the spread of misinformation about COVID-19. The authors collect a corpus of misinformation stories and use Latent Dirichlet Allocation (LDA) to reveal latent narratives around a variety of topics. The extracted topics led the authors to associate online conversation with conspiracy theories. A similar method for identifying shifts in a narrative, is extracting causal relationships that exist between narrative events. Causal relation extraction and arrangement tasks are popularly used for story completion. Yusuke *et al.* [100] propose a story comprehension task called “Missing Position Prediction” (MPP), where

<sup>3</sup><https://www ldc.upenn.edu/collaborations/past-projects/ace>

<sup>4</sup><https://framenet.icsi.berkeley.edu/fndrupal/>

the objective is to predict the position of a missing narrative component given an incomplete story as input. The authors use the ROCStories dataset [101], which is an example of one of many such datasets that contain non-fictional daily-life occurrences and stories. Datasets such as *ROCStories* and *NewsWire*<sup>5</sup> are commonly used as sources for Natural Language Understanding tasks, such as recognizing causal relationships between events. Zhichao et al. [57] claim that the source *NewsWire* can be augmented with additional data sources such as scripts and social media blogs, to learn fine-grained casual relations that may not be found in traditional causal knowledge bases learned from *NewsWire*.

The method of sequencing casual events and related shifts is known as *storyline construction and evolution*. A similar line of work is *Timeline Generation*, the process of extracting events and placing them on a chronological timeline, based on a user queries from a collection of documents [22], [87], [153]. There are several works done on transforming event structures into storylines and timelines. For example, Croft et al. [26] propose an event decomposition method that expresses each participant in a causal event as individual temporally subevents, which allows the authors to represent storylines as an evolving set of subevent interactions over time. The authors also apply the event decomposition approach to event annotation for storyline analysis.

Storylines can also formally be organized as *story chains*, which are essentially sequential events that together form a story or narrative. Zhu et al. [159] define a story chain as “a construction of news articles that reveal hidden relationships among different events”. The authors create an algorithm that uses random walks on a bipartite graph and keyword-based search to form a coherent and accurate story chain based on a user’s query. Prior to this particular work, several research projects have ordered news articles based on hierarchical structures [42], [99], [104], [130]. Structurally organizing sequential events, whether in a graphical or timeline form, is a useful tool for efficiently conveying chronological summaries of large and fragmented narratives. Evolutionary Timeline Summarization (ETS) is the process of producing evolutionary summaries of narratives related to general news queries. Typically, an *evolution trajectory* is returned in a timeline structure, with correlated summaries at important dates [153]. Pasquali et al. [115] develop an online tool based on a keyphrase extraction algorithm, that allows users to generate narrative summarizations for user search queries. Another timeline summarization tool developed by McCreddie et al. [98] automatically extracts events from social media websites over time, and issues timeline updates to users.

### C. NARRATIVE GENERATION

The correlative component of narrative construction, is the process of automatically building coherent and fluent story passages, generally falling under the subfield of

*text generation*. Traditional text generation applications include neural machine translation [8], question-answering [155], and summarization [105]. Generation of text is useful for tasks that involve event or storyline prediction. There are generally three broad methods for generating narratives (structural, goal-oriented, transformer model-based), as shown in Figure 2.

One example of a structural-based method is a hierarchical approach to neural story generation. Fan et al. [36], use semi-supervised methods to generate coherent text, based on a predicted story premise. Given a large dataset of human written stories and their paired writing prompts, a sequence-to-sequence model is trained to generate an initial premise, then structurally transforms the premise into a coherent passage of text. Similarly, Martin et al. [96] use sequential learning, where textual story data is preprocessed into event sequences, and later decomposed into two sub-problems of event generation from event representations (event2event), and sentence generation from event representations (event2sentence).

A progression from sequence-to-sequence models is the use of large language models like GPT [119], Grover [157], and CTRL [71], to generate fluent and linguistically accurate text. For example, GPT-2 has been used to generate texts in a variety of domains like law [80], cybersecurity [120], and literature [34]. GPT-2 has even been used to continue dialogue given contextual online social media conversations, such as those found on Reddit [158]. Though neural language models have shown great promise in generating text with grammatical and linguistic consistency, there are known issues in using these models to generate text with logical coherency, opening a broad area of research that involves logical improvement of generated text [69], [74]. Examples of goal-oriented methods for generating stories involve tasks that aim to improve the generated text in stylistic, event-driven, or context-driven scopes. Li et al. [86] propose a framework that first maps text to *VerbNet* frames and predicates. The task is for the GPT-2 model to generate components of a story based on sets of most probable sentence candidates, providing the ability for GPT-2 to increase the logical coherency of generated text. Similarly, Mao et al. [94], improve automatic story generation by using a two-stage GPT-2 fine-tuning process, to augmenting traditional neural language models, with the ability to generate text using common sense grounding methods.

### D. EVALUATING EXTRACTED AND GENERATED NARRATIVES

Once narratives are either extracted or generated, it is an important consideration to evaluate both the narrative structure, as well as, factors like bias, factual correctness, and the persuasion factors that may influence a sequence of events.

#### 1) BASELINES FOR NARRATIVE UNDERSTANDING TASKS

There is a general lack of baseline datasets available for evaluating both constructed and generated narratives.

<sup>5</sup><https://www.newswire.com/>

However, there are a few benchmarks that are widely used in the narrative understanding community.

One example of a standard and commonly used benchmark for evaluating generic story understanding is the *Story Cloze* task [102], where the goal is to learn commonsense narrative sequences. This benchmark is one of the few defined benchmarks that exist for evaluating a series of narrative tasks, such as story completion [133], causal relation extraction and arrangement [131], and story generation [94]. Story Cloze is inspired by a predecessor benchmark called *Narrative Cloze* [18], which optimizes single event prediction, given a sequence of events. It was the first known published narrative benchmark, and grew from two traditional lines of research in summarization (topic networks) and anaphora resolution (case frame networks). Narrative Cloze, stems from the *Cloze Task* which evaluates a system (or human being), by randomly removing a word from an input sentence, and testing the ability of the system to fill in the blank [142]. Similarly, Narrative Cloze is a sequence of narrative events, where one event is randomly removed. The task tests the ability of the system to predict the missing verb and typed dependency. The task operates on the Newswire dataset and aims to learn narrative relations between co-referring events, partially orders temporally connected events, and prunes self-contained narrative chains from a space of existing events. Story Cloze extends Narrative Cloze to evaluate tasks that learn concluding *sentences* to an input story, rather than limiting identification to single events.

Linking Models of Lexical, Sentential and Discourse-level Semantics (LSDSem)' 17 used the benchmark as a shared task (Story Cloze Test), where the goal was to predict the correct ending of a story given a four-sentence story and two possible endings [102]. The competing models and results for the task are provided in Figure 3. It is evident that diverse combinations of models, pre-trained resources, and tools can be used to optimize story completion tasks. Though these results are from the 2017 shared task, Story Cloze is has consistently been adapted using a variety of methods such as graph models [106] and more modern large language models such as BERT [152].

The tasks above mainly address story completion tasks, given potential endings in self-contained narratives. Zhu *et al.* [159] describe an example of a system that tackles the research problem of story completion for fragmented narratives across a variety of news articles. An example of the story chain output is provided in Figure 4.

First, a set of chronologically ordered articles is created based on keyword searches. Though the articles come from various fragmented sources, this study is limited to filling in story components given start and end articles. The authors define the fragmented story chain problem as a divide and conquer search problem. Given an initial story chain with start and end articles, the algorithm iteratively performs two tasks: (1) search and retrieval of relevant articles and (2) pruning irrelevant and redundant articles, until there are no remaining articles in the chronologically ordered set. To rank the

relevancy of articles, the authors use random walks on a bipartite graph and add nodes with the highest scores to the story link. The final story chain contains multiple links, allowing users to visualize fragmented perspectives for a given topic. An example of one story chain given a query and start and end articles is shown in Figure 4.

It is important to note that the start and end articles inherently guide the construction of the rest of the story chain, adding bias to the constructed stories. Apart from this specific system, combating bias is generally an ongoing limitation for constructing and generating story chains. The rest of this section describes examples of evaluations used to measure bias in computational narrative tasks.

As shown under the evaluation section in Figure 2, there are several existing methods that are used to measure bias, correctness of information, and potential persuasion techniques that may influence the construction and generation of narratives. The notion of *bias* has become an important evaluation factor for both the ethical and fair generation and construction of narratives. In terms of generation, several research works have addressed questions surrounding associated contextual biases generative models have when forming narratives. For example, Lucy *et al.* [91] found that stories generated by the latest GPT-3 model, may exhibit gender stereotypes. The authors remove character names from the generated GPT-3 text and train a topic model to derive the top 50 topics in any given text and calculate the probability of a topic occurring with feminine or masculine characters. They found that feminine characters are more likely to be discussed surrounding topics like family and emotions, while male characters surround topics like politics and sports. Similarly, Magee *et al.* [93] conduct a study to examine the observation of intersectional bias in GPT-2 and GPT-NEO text generation across three combined social categories of gender, religion, and disability. The authors found that bias exists at significant levels across three social categories. Controlled, or targeted text generation, can also help in eliminating biases. For example, Dathathri *et al.* include a language detoxification layer, trained on a toxicity dataset before generating text, to eliminate potential biases in the input prompts. Understanding generated biases from large language models like GPT-3, can help us avoid misconstrued and incorrect story chains, or even narratives generated from singular perspectives [28].

In terms of eliminating bias when constructing narratives (particularly, fragmented narratives), traditional research in bias identification and mitigation can be employed when chaining sources together into a constructed storyline. There have been several research works conducted to first identifying bias, as well as mitigating bias. In recent years, researchers have begun to report potential biased metrics across both language models and datasets [24], [44]. In addition, other researchers have created standards for evaluating model fairness [92], [123]. Transformer-based language models and multidimensional word embeddings have

Rank	CodaLab Id	Model	ROCStories	Pre-trained Embeddings	Other Resources	Accuracy
1	<b>msap</b>	Logistic regression	Spring 2016, Winter 2017	—	NLTK Tokenizer, Spacy POS tagger	<b>0.752</b>
2	<b>cogcomp</b>	Logistic regression	Spring 2016, Winter 2017	Word2Vec	UIUC NLP pipeline, FrameNet, two sentiment lexicons	0.744
3	<b>tbmihaylov</b>	LSTM	—	Word2Vec	—	0.728
4	<b>ukp</b>	BiLSTM	Spring 2016, Winter 2017	GloVe	Stanford CoreNLP, DKPro TC	0.717
5	<b>acoli</b>	SVM	—	GloVe, Word2Vec	—	0.700
6	<b>roemmele</b>	RNN	Spring 2016, Winter 2017	Skip-Thought	—	0.672
7	<b>mflor</b>	Rule-based	—	—	VADER sentiment lexicon, Gigaword corpus PMI scores	0.621
8	<b>Pranav.Goel</b>	Logistic regression	Spring 2016, Winter 2017	Word2Vec	VADER sentiment lexicon, SICK data set	0.604
9	<b>ROC NLP (baseline)</b>	DSSM	Spring 2016, Winter 2017	—	—	0.595

FIGURE 3. Models, resources, and results for teams participating in the 2017 story cloze shared task [102].

Query intent	Start $d_s$ and end $d_t$ articles
C1: How Hurricane Katrina is related to government policies	$d_s$ : Hurricane Katrina hit New Orleans (08/26/05) $d_t$ : Attacking Bush, Clinton Urges Government Overhaul (04/14/07)
C2: How Japan earthquake has impact on nuclear policy in German nuclear company	$d_s$ : Japan super quake, tsunami terrify tremor-prone nation (03/11/11) $d_t$ : E.ON to sue German government over nuclear closure (11/14/11)
C3: How Japan earthquake has impact on competition between Toyota and Volkswagen	$d_s$ : Japan super quake, tsunami terrify tremor-prone nation (03/11/11) $d_t$ : Volkswagen may topple Toyota as world's top automaker (10/24/11)
C4: Story about O.J. Simpson trial	$d_s$ : O.J. Simpson's Ex-Wife Found Dead in Double Homicide (06/13/94) $d_t$ : Simpson jury reaches verdict in six hours (10/02/95)
C5: How O.J.Simpson trial has impact on racial problems	$d_s$ : O.J. Simpson's Ex-Wife Found Dead in Double Homicide (06/13/94) $d_t$ : Race flares anew as polarizing issue in U.S. life (10/19/95)

#### Algorithm 1 Story chain finding algorithm

Input: chronologically-ordered articles  $d_1, d_2, \dots, d_n$ , Start node  $s$ , End node  $t$   
Initialize story chain  $C = s - t$ , input link  $l = \{s - t\}$   
**repeat**  
1. Pruning process (a): Prune least relevant articles  
2. Select a best article  $a_i$ , that can be added to the link. Story chain becomes  $C = s - a_i - t$ .  
3. Pruning process (b): Prune redundant articles  
4. Update input link as  $l = \{s - a_i, a_i - t\}$ . Repeat step 1, 2 and 3 for each of the input link in  $l$   
**until** There are no articles left in the set. (Articles are either been added to the chain or have been pruned.)  
Output: story chain  $C = s - a_1 - a_2 - \dots - a_i - t$

- 1) A Blast of Rain but Little Damage as Hurricane Hits South Florida 2005 8 26
- 3) Hurricane Drenches Florida And Leaves Seven Dead 2005 8 27
- 4) FEMA, Slow to the Rescue, Now Stumbles in Aid Effort 2005 9 17
- 5) Millions Are Still Without Power and in Need of Basic Supplies 2005 10 26
- 6) South Florida Scrambling To Find Emergency Housing 2005 11 11
- 7) Bush Erred In Responding To Katrina, Lamont Says 2006 8 25
- 8) Bush failed to act immediately after Hurricane Katrina to waive the requirement that state and local governments match federal rebuilding funds 2007 2 13
- 9) Bush Consoles Victims of Tornadoes in Alabama and Georgia 2007 3 04
- 10) Gulf Hits Snags In Rebuilding Public Works 2007 3 31
- 11) Attacking Bush, Clinton Urges Government Overhaul 2007 4 14

FIGURE 4. Example of story completion task for fragmented articles. [159].

been used to detect bias via cosine similarity of individual tokens [52], [65], [84], [97], [139]. Another popular method for detecting bias is through sentiment analysis approaches [62], [68], [143]. Sentiment analysis has been found to provide a more broad measure for biases that do not directly relate to contextual meanings of words such as

genders or roles, but rather the intersection of many social influences [93].

Despite the uses of computational methods to eliminate bias across generated stories, it is also valuable to create methods to eliminate bias across human annotators who often evaluate generated or constructed narratives.



Nguyen *et al.* [107] study the effects of idea and story generation within groups vs individuals. Their study found that groups naturally eliminate personal bias by offering multiple interpretations of events. As a result, the groups were able to generate more comprehensive, complete, and unique stories, in comparison to those created by singular individuals. Though there is significant effort in eliminating unintended bias across generating narratives, there are some narratives that are purposefully crafted to reach certain individuals or groups of individuals. There are several instances across social media where the use of online persuasion has caused a rise in targeted narratives across specific social groups, such as political parties. This phenomenon has been shown to cause group-think, and echo-chamber like behavior across various online communities [3].

The most traditional form of using persuasion to target specific communities is through advertisements [29]. For example, Huang *et al.* [58] expand the concept of *narrative transportation* [35] to branded storytelling in social media, and study online brand advertisements through a narrative persuasion perspective. Similarly, Dey *et al.* demonstrate the role of crowdfunding campaigns in influencing individual opinion on stigmatized topics [30]. Brand to consumer relationship building such as these, fall under the categories of *crowdsourcing and crowdfunding* techniques, which allow brands to seek opinions, knowledge, ideas, and resources from groups of consumers. In this scenario, brands are purposefully allowing consumers to influence narratives surrounding their products. The crowdsourcing and echo chamber phenomena opens an interesting research area when considering user-centric constructed or generated narratives. A similar phenomenon is observed when considering bias impacts on narratives surrounding news topics such as, social movements [3]. A possible future evaluation could be verifying the sources (either platforms or users) of narrative components, ensuring there is no group bias when training generative models, and eliminating an inconclusive chaining of fragmented sources to form a comprehensive narrative.

In addition to eliminating bias, *fact-checking* is also a common procedure when evaluating generated and constructed narratives. When considering online narratives, fact-checking of the input text is especially crucial due to the large presence of *bots*, which have the potential of influencing online discourse and narrative development [23], [118]. Bots are often also part of campaigns that aid in spreading *misinformation* [73]. There are several methods proposed to evaluate the factual correctness of a given text. Hassan *et al.* developed *ClaimBuster*, a system that uses a combination of natural language processing and database query techniques to aid in the process of fact-checking [54]. The system monitors discourse across social media and online news sources, and matches factual claims to a repository of fact checks from human experts. Similarly, Shao *et al.* develop *Hoaxy*, an extensible framework for locating online

misinformation, and checking text correctness against fact-checking databases [132].

### E. OPEN RESOURCES

There are several existing tools, ontologies, and datasets that provide the ability to consistently extend, evaluate, and train various types narrative understanding systems discussed in this survey. An overview of the existing resources is displayed in Table 3. The *tools* category describes software-based systems that can be leveraged by researchers to aid in annotation, visualization, as well as data aggregation. One example of a tool available to aid in narrative annotation tasks is Brat2Viz [4]. The authors extend the standard BRAT rapid annotation tool [140] to graphically represent narrative components (actors, events, and times) as well as the semantic relationships between the them. The visualization can aid human analysts in maintaining ongoing, developing plots, as well as, seek potentially new directions in existing plots. Similar to annotation tools like Brat2Viz, there are other proposed research works that aim to visually represent components of a narrative in a user interface. One such example is by Segel *et al.* who use hierarchical visualizations of online data such as news articles, budget forecasts, and employment rates [128].

Apart from visual representation, there is a larger research area that concentrates efforts on representing narratives through formal structures such as data models and ontologies, with the goal of extending with existing pipelines that leverage automated narrative analysis. Unlike ontologies, data models do not represent generic domain semantics but instead, represent schemas for fine-grained tasks. The goal of a narrative-based data model is to provide a schema to represent the format, structure and compositional features of narratives [76]. Akimoto *et al.* formalize a hierarchical framework that integrates chronological organization of entities, narrative events, and discourse within literary based narratives [2]. Similar narrative representation schemas have additionally been developed through language processing techniques, like distributed word embeddings. Lee *et al.* describe methods to create *story representations* to represent relationships between characters, as well as their social roles and thematic story elements [81].

Data models can be integrated with general knowledge ontologies for a diversity of applications and subtasks. Tuffield *et al.* propose methods for extending a taxonomy of characters, plots, and user-based narrative features into an ontological format that represents chronological ordering of events, the stories that are fulfilled by the connection of the events, and the general narrative type [146]. Often, extensibility for narrative representations can be achieved through ontologies that model aspects of events themselves. For example, Segers *et al.* describe the Circumstantial Event Ontology (CEO), that builds upon an existing event ontology called the Event and Implied Situation Ontology (ESO) for capturing chains of calamity events in Newswire [129].



Similar event ontologies include Linking Open Descriptions of Events (LODE),<sup>6</sup> Simple Event Model (SEM),<sup>7</sup> Rich Event Ontology (REO) [13] and Comprehensive Event Ontology (CEVO) [134]. One example of extending event classes from event data models is exemplified by Shekarpour *et al.*, where the authors describe methods to create sub-event classes for specific news related event categories with the goal of real-time extraction of events across fragmented news sources [135].

In addition to data models and ontologies, training datasets can be used to train AI-based systems to recognize narrative elements. Several research efforts have developed narrative datasets for construction and generation-based tasks, as well as for benchmarking purposes. As discussed in Section III-B, events and sequences of events are one of the primary components that form all narratives. One example of a widely used benchmark for Event Detection and Tracking (EDT), Relation Detection and Characterization (RDC), and Event Detection and Characterization (EDC) tasks is known as Automatic Content Extraction (ACE) [31]. This dataset is widely used in the narrative understanding community, for subtasks related to events [88].

*StoryDB* is another benchmark example that focuses on narrative understanding tasks. It is a large multilingual corpus that contains about 21,000 stories in 42 different languages [144]. Examples of applications that can leverage *StoryDB* include cross-cultural research for narrative structure, classification of narrative types, summarization of a narrative, and end-to-end narrative generation. Though *StoryDB* is an example of a large and general corpus of text that represents stories across various languages, cultures, and topics, there also exist niche datasets that concentrate efforts on sub-fields within the narrative understanding domain. One such example is the Literature Summary and Character Understanding (LiSCU) dataset, which is a corpus of literary summaries paired with descriptions of characters appearing in the summaries [12]. This dataset can be used for downstream tasks that involve forming opinions on characters throughout a narrative. The authors test the dataset with a character description generation task, to evaluate the character-centric understanding of narratives. Additional niche narrative datasets include *CompRes*, a dataset for narrative structure in news [83], and *CONAN*, a multilingual text dataset of hate speech and counter-narrative pairs [37]. In addition to textual narrative datasets, there is also an emergence of visual storytelling datasets such as *Visual Story Dataset (VIST)*, that models unique photos in narrative sequences, aligned to descriptions of the photos in sequence, as well as individually [59].

## V. FUTURE DEVELOPMENTS

Thus far, we have reviewed definitions and types of different narratives, primary components of universal narratives,

as well as several research areas that can be used to construct, generate, and evaluate narratives. These research areas include many open challenges, and potential for ongoing improvements. As shown in Table 3, there are a multitude of limitations associated with each current research area. At a general level, primary limitations surround gaps in generalizability, logical coherence, implementation of bias measures, and resource constraints. This section will briefly review the primary limitations for each research area and present potential future tasks that could lead to continuing progress.

### A. NARRATIVE EXTRACTION

Given the many types of narratives, as well as diversity of potential entities and semantic relationships that compose differing narratives, in Section IV-A, we have seen that a number of methods used for narrative extraction are typically based on targeted inputs for narrative-type specific tasks. The major limitation for narrative extraction areas is the challenge for generalizability. Many of these methods, whether concerning detection and extraction of events and semantic relationships, or creating narrative representation languages, are based on targeted inputs, and cannot be applied to multiple domains easily. For example, the narrative representation language for a fantasy novel will likely not be backwards compatible to a representation language of a strategic news-based narrative. Similarly, events and semantic relationships in free-text can differ tremendously depending on the domain. More discussion on domain-specific studies for event detection and relationship extraction is described more thoroughly in Section IV-A. However, creating domain-specific datasets for novel narrative problems can be labor intensive and costly. A future research area can work towards creating generalizable representations that fit multiple narrative types.

One of the greatest challenges in narrative understanding tasks is integrating domain-agnostic, commonsense knowledge for global interpretation of a variety of narrative events. Generalizability can produce many benefits, such as decreasing human intensive narrative extraction tasks, and providing reference models for story chain-based information flow.

Achieving generalizability is expected for innovation across the major narrative understanding tasks discussed in this survey. Both general and domain-specific texts can benefit from models that have the ability to automatically and accurately chain events into event chains. Narratives that concern colloquial events such as news stories, and personal occurrences, can more easily generalize in comparison to domain-specific events due to the sparsity of available domain-specific narrative training. In addition, there is no current research in developing cross-domain methods for extracting narrative components, apart from methods that chain and generate stories based on general domain narratives samples derived from datasets like ROCStories [101]. More information on ROCStories can be found in Section IV-B.

<sup>6</sup><https://linkedevents.org/ontology/>

<sup>7</sup><https://semanticweb.cs.vu.nl/2009/11/sem/>

**TABLE 3. Computational understanding of narratives: Current research and open challenges.**

Category	Focus/Objective	Contributions	Limitations
<b>Narrative Extraction</b>			
<b>Event Detection</b> [20], [27], [122], [138]	Identification, tracking, and characterization of event entity types in unstructured and semi-unstructured text data.	Describes approaches to categorize events into types, sub-types through methods such as hierarchical clustering, semantic frame parsing, and large language modeling.	<ul style="list-style-type: none"> <li>Lacks statistical abilities to detect events defined implicitly in text.</li> <li>Performance of methods rely on quality of designed input features, and unable to generalize well.</li> </ul>
<b>Relationship Extraction</b> [10], [109], [127]	Extraction of semantic relationships that exist between named entities.	Methods to link narrative events together based on contextual relationships between entities.	Tasks tested on small scale typically limited to domain-specific applications.
<b>Narrative Representation Language</b> [46], [60], [154]	Representation of primary components of a narrative through extensible, machine-readable narrative representation structures.	<ul style="list-style-type: none"> <li>Creation of functional story schemas that capture latent, structural features in unstructured text.</li> <li>Visualization techniques to display fragmented narrative components in a user interface.</li> <li>Development of domain-specific narrative schemas.</li> </ul>	<ul style="list-style-type: none"> <li>Subject bias when forming grammar rules.</li> <li>Non-generalizable representations.</li> <li>Schema clustering algorithms easily lose ordering/temporal information.</li> </ul>
<b>Story Evolution</b>			
<b>Shift Detection</b> [61], [95]	Identifies transitions between plot elements.	<ul style="list-style-type: none"> <li>Methods to study shifts in narratives across disparate online sources.</li> <li>Topic Modeling to assess shifts in narrative sentiments</li> <li>Extraction of chronological themes across social media narratives.</li> </ul>	<ul style="list-style-type: none"> <li>Non-generalizable models.</li> <li>Scalability issues.</li> <li>Temporal expression limited to dates and timestamps.</li> </ul>
<b>Storyline Construction</b> [26], [42], [99], [104], [130], [159]	Create methods for sequencing causal events and temporal narrative shifts.	<ul style="list-style-type: none"> <li>Event decomposition methods to represent storylines as evolving sets of subevent interactions.</li> <li>Formation of story chains.</li> <li>Ordering of online sources based on hierarchical structures.</li> </ul>	<ul style="list-style-type: none"> <li>Targeted inputs with known start and end to storyline.</li> <li>Does not study the potentials of reordering of events based on new information.</li> <li>Limited analysis on unrealized events in annotations and representations.</li> </ul>
<b>Timeline Generation and Summarization</b> [22], [87], [98], [115], [153]	Extracting events and placing them on a chronological timeline, based on a user queries from a collection of documents.	<ul style="list-style-type: none"> <li>Methods to structurally organize sequential events in both graphical and timeline formats.</li> <li>Producing evolutionary summaries (evolution trajectory) of events.</li> </ul>	<ul style="list-style-type: none"> <li>Does not take into consideration aspects such as bias and hyper-partisan augmentation.</li> <li>Most systems developed worked independently of an integrated search engine.</li> </ul>
<b>Causal Relation Extraction and Arrangement</b> [57], [100]	Identifying cause/effect relationships between narrative events.	<ul style="list-style-type: none"> <li>Methods that predict position of missing narrative components, given incomplete story inputs.</li> <li>Augmentation of current baseline datasets to learn additional fine-grained causal relations.</li> </ul>	<ul style="list-style-type: none"> <li>Experimentation on short stories, minimizing contextual feature space.</li> <li>Domain-specific experiments not generalizable.</li> </ul>
<b>Narrative Generation</b>			
<b>Structural Methods</b> [36], [96]	Preprocessing narrative text into structures such as sequences or hierarchical graphs and later decomposed into task-oriented text generation.	<ul style="list-style-type: none"> <li>Generation of linguistically coherent text given an incomplete narrative structure as input.</li> <li>Using sequential learning to preprocess textual story data into event sequences.</li> <li>Methods to reduce sparsity of events to maintain semantic meaning of textual story data.</li> </ul>	<ul style="list-style-type: none"> <li>Introduces sub-problem of potentially unreadable output.</li> <li>Random sampling can occasionally miss important tokens.</li> <li>Sequence modeling introduces content repetition as output is generated.</li> </ul>
<b>Goal-oriented Methods</b> [86], [94]	Improving generated text in logical, stylistic, event-driven, and context-driven scopes.	<ul style="list-style-type: none"> <li>Editing beam search mechanisms to generate likely next sentence based on logical context.</li> <li>Increased logical coherency of generated text.</li> </ul>	<ul style="list-style-type: none"> <li>Lack of characterization of narrative actors in generation pipelines.</li> <li>Frame Parsers fail at processing complex sequences.</li> </ul>

**TABLE 3.** (Continued.) Computational understanding of narratives: Current research and open challenges.

			<ul style="list-style-type: none"> <li>• Reliance on short-term common sense grounding poses limitations on achieving story-level coherence.</li> </ul>
<b>Large Language Modeling</b> [34], [158]	Using large language models to generate fluent and linguistically accurate narratives.	<ul style="list-style-type: none"> <li>• Automatic generation of large samples of domain-agnostic and linguistically coherent text.</li> <li>• The ability to generate plausible narratives in a variety of domains through the use of fine-tuning.</li> </ul>	<ul style="list-style-type: none"> <li>• Lack of logical coherence in generated text.</li> <li>• Biases often present in generated output.</li> </ul>
<b>Eliminating Bias</b> [28], [91], [93]	Detecting contextual biases from generated and constructed narratives.	<ul style="list-style-type: none"> <li>• Comprehensive analyzation of potential biases in generated output from large language models.</li> <li>• Identification of bias in popular models.</li> <li>• Methods in combating bias.</li> </ul>	<ul style="list-style-type: none"> <li>• Prompt selection for bias tests can also include levels of bias themselves.</li> <li>• Efficacy of fine-tuning biases have limited evaluations.</li> <li>• Lack of diversity in datasets used to train bias detection classifiers.</li> </ul>
<b>Fact Checking</b> [54], [132]	Methods to evaluate factual correctness of a given text.	<ul style="list-style-type: none"> <li>• Systems that integrate natural language models and fact checking databases to evaluate textual correctness.</li> <li>• Extensible frameworks for flagging potential misinformation.</li> </ul>	<ul style="list-style-type: none"> <li>• Lack of baselines to evaluate misinformation detected.</li> <li>• System not integrated with general purpose search engines, but rather with small sets of online sources.</li> </ul>
<b>Bot Influence</b> [23], [73], [118]	Study impacts of bots influencing narrative campaigns and online influence.	Evidence of bots used for spreading misinformation campaigns and narratives.	Major focus of research on detection of bots, but limited work on influence features that motivate bots.
<b>User Profiling and User Behavior Modeling</b> [29], [58], [107]	Eliminating bias across human annotators during evaluation of generated or constructed narratives.	<ul style="list-style-type: none"> <li>• Extensive studies of group-bias and echo-chamber impacts in story generation tasks.</li> <li>• Studies of attributes of narratives targeted towards individuals or direct groups of people.</li> </ul>	<ul style="list-style-type: none"> <li>• Does not consider repeatability factors for each study.</li> <li>• Lack of pruning methods to modify data to match respective societal norms.</li> </ul>
<b>Open Resources</b>			
<b>Datasets</b> [12], [31], [37], [59], [83], [144]	Creation of diverse narrative datasets that can be used as training and baseline resources for construction and generation tasks.	<ul style="list-style-type: none"> <li>• Multilingual narrative representation.</li> <li>• Domain-specific narrative datasets.</li> <li>• Combined visual and text feature narrative dataset.</li> </ul>	Lack of community-wide baseline (standards only exist for traditional related tasks, such as event detection).
<b>Tools</b> [4], [128]	Development of software-based systems that aid in annotation, and visualization of constructed narratives.	User Interfaces to easily update, extend, and chronologically related events.	High concentration on visualization, with limited work done to integrate methods that improve construction-based tasks.
<b>Ontologies and Data Models</b> [2], [135], [146]	Representation of narrative structures in ontological formats.	Integration of narrative elements such as actors, events, and plots into new and existing extensible ontologies.	<ul style="list-style-type: none"> <li>• Difficult to represent nested structures in online discourse and traditional story formats.</li> <li>• Only static narrative features have interfaces in existing narrative ontologies.</li> <li>• Temporal features not integrated.</li> </ul>

One approach to achieving cross-domain narrative generalizability is the development of intermediate tasks that map domain-specific components to general narrative schemata. In this way, we can potentially apply general narrative tasks to domain-specific texts. One approach is integrating general plot-based elements (Table 2) across subject matters, allowing for narrative schemata domain transfer learning capabilities. For example, the vocabulary used to describe the climax of a fantasy novel will likely be different than the climax

vocabulary of an article describing the severity of a virus variant. However, using the contextual meaning of the climax (occurrence of the major event of the story) can help identify and compare examples of plot elements across different domains. To achieve this, we can first gather examples for each plot element across multiple domains and apply them as training data for cross-domain schema learning tasks. The primary associated task involves investigating methods that allow unsupervised natural language models to recognize

similarities and differences across examples for each plot element and evaluating the model on unseen domains [110].

### B. STORY EVOLUTION

As discussed previously in Section IV-B, creating story-chains, or sequences of events that constitute a given narrative, are largely based on temporal features such as time-stamps and dates. Though these methods have proven to be useful for tasks that involve a strict layout of events that lead from a starting point to a particular outcome (pre-determined plots, or factual sets of sequenced events). However, when considering research areas that aim to locate real-time, developing narratives across fragmented sources, event ordering does not necessarily only develop chronologically, but also can be influenced by socialization factors such as peaks in online conversation surrounding specific topics, updated information and additional evidences, as well as user bias. Future research areas can potentially explore such socialization factors when constructing the ordering of events.

### C. NARRATIVE GENERATION

The lack of logical coherence in generated text presents a large research gap in text and narrative generation tasks. This leads to methods such as structural and goal-oriented text generation, which have additional limitations such as difficulty in handling complex sequences, unreadable output, and content repetition. While state-of-the-art large language models are better at handling complex text and can generate linguistically accurate text, the central limitation is the lack of logical reasoning during generation. There are several ongoing research projects that explore grounding state-of-the-art language models with knowledge representations to output targeted text [48], [63], [148].

### D. EVALUATION

Though there are several narrative construction and generation methods already developed and in-progress, there is a lack of consistent baselines used to measure the accuracy, fluency, and reputability of such methods. Primary limitations surrounding evaluations include low variety of community-wide baselines and increasing gaps in studying bias and profiling when constructing or generating narratives. Future work can include design, testing, and construction of training datasets and baselines for both text generation and construction-based tasks. Research to include bias measures when creating training datasets and baselines is also an ongoing and essential research task.

### E. OPEN RESOURCES

As discussed previously, many of the narrative-based ontologies, data models, and datasets available openly are domain-specific to particular types of narratives. Developing methods toward generalizing the representation of narrative structures is a large research opportunity. In addition to creating more datasets, data models, and ontologies, there

are also numerous opportunities for creating additional data visualization tools to illustrate the components of narratives and temporal aspects of narrative events to human analysts. Though the use of computational methods can augment narrative analysis, the goal is not to remove human analysts from the pipeline and replace them with computational systems. Rather, a large goal is to use computational systems to better represent and provide consumable and relevant data for human analysts.

## VI. CONCLUSION

The novel contribution of this paper is the collation and documentation of narrative understanding methods and challenges associated with narrative construction, generation, representation, and evaluation tasks, as well as universal definitions for different types of narratives, and the common components that compose them. There is no other resource to our knowledge that provides an overview the above topics in formal documentation. While we have reviewed the narrative understanding literature extensively and provided a hierarchy of primary research areas (Figure 2), there are still some topics that have not been addressed by the current literature and will require community-wide research, engineering collaboration to address. These challenges are documented in Table 3 and Section V. Our goal is to provide the narrative understanding community with this document to archetype current research areas, challenges, and potential avenues for innovation.

## REFERENCES

- [1] S. Abiteboul, P. Buneman, and D. Suciu, *Data on the Web: From Relations to Semistructured Data and XML*. San Mateo, CA, USA: Morgan Kaufmann, 2000.
- [2] T. Akimoto, "Computational modeling of narrative structure: A hierarchical graph model for multidimensional narrative structure," *Int. J. Comput. Linguistics Res.*, vol. 8, no. 3, pp. 92–108, 2017.
- [3] E. Amenta, T. A. Elliott, N. Shortt, A. C. Tierney, D. Türkoglu, and B. Vann, "From bias to coverage: What explains how news organizations treat social movements," *Sociol. Compass*, vol. 11, no. 3, Mar. 2017, Art. no. e12460.
- [4] E. Amorim, A. Ribeiro, I. Cantante, A. Jorge, B. Santana, S. Nunes, M. D. P. A. Silvano, A. Leal, and R. Campos, "Brat2viz: A tool and pipeline for visualizing narratives from annotated texts," in *Proc. Text2Story Workshop Narrative Extraction Texts, Eur. Conf. Inf. Retr. (ECIR)*, 2021, pp. 1–8.
- [5] Global Security Initiative Arizona State University. (2021). *Narrative, Disinformation and Strategic Influence*. [Online]. Available: <https://globalsecurity.asu.edu/expertise/narrative-disinformation-and-strategic-influence>
- [6] B. Arnold and B. Eddy, *Visual Storytelling*. New York, NY, USA: Thomas Delmar Learning, 2007.
- [7] N. Bach and S. Badaskar, "A review of relation extraction," *Literature Rev. Lang. Statist. II*, vol. 2, pp. 1–15, Nov. 2007.
- [8] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," 2014, *arXiv:1409.0473*.
- [9] M. A. Bock, "Newspaper journalism and video: Motion, sound, and new narratives," *New Media Soc.*, vol. 14, no. 4, pp. 600–616, Jun. 2012.
- [10] A. Bodrova and V. Bocharov, "Relationship extraction from literary fiction," in *Proc. Dialogue, Int. Conf. Comput. Linguistics*, 2014, pp. 1–8.
- [11] C. Bosangit, S. McCabe, and S. Hibbert, "What is told in travel blogs? Exploring travel blogs for consumer narrative analysis," in *Proc. Int. Conf. Inf. Commun. Technol. Tourism*, Amsterdam, The Netherlands. Vienna, Austria: Springer, 2009, pp. 61–71.



- [12] F. Brahman, M. Huang, O. Tafjord, C. Zhao, M. Sachan, and S. Chaturvedi, "Let your characters tell their Story: A dataset for character-centric narrative understanding," 2021, *arXiv:2109.05438*.
- [13] S. Brown, C. Bonial, L. Obrst, and M. Palmer, "The rich event ontology," in *Proc. Events Stories News Workshop*, 2017, pp. 87–97.
- [14] R. Campos, A. Jorge, A. Jatowt, and S. Bhatia, "The 3rd international workshop on narrative extraction from texts: Text2story 2020," in *Proc. Eur. Conf. Inf. Retr.* Cham, Switzerland: Springer, 2020, pp. 648–653.
- [15] R. Campos, A. Jorge, A. Jatowt, S. Bhatia, and M. Finlayson, "The 4th international workshop on narrative extraction from texts: Text2story 2021," in *Proc. Eur. Conf. Inf. Retr.* Cham, Switzerland: Springer, 2021, pp. 701–704.
- [16] J. Carroll, "The deep structure of literary representations," *Evol. Hum. Behav.*, vol. 20, no. 3, pp. 159–173, May 1999.
- [17] R. Chaffin and D. J. Herrmann, "The similarity and diversity of semantic relations," *Memory Cognit.*, vol. 12, no. 2, pp. 134–141, Mar. 1984.
- [18] N. Chambers and D. Jurafsky, "Unsupervised learning of narrative event chains," in *Proc. ACL, HLT*, Jun. 2008, pp. 789–797.
- [19] R. Chen and Y. Sakamoto, "Perspective matters: Sharing of crisis information in social media," in *Proc. 46th Hawaii Int. Conf. Syst. Sci.*, Jan. 2013, pp. 2033–2041.
- [20] Y. Chen, L. Xu, K. Liu, D. Zeng, and J. Zhao, "Event extraction via dynamic multi-pooling convolutional neural networks," in *Proc. 53rd Annu. Meeting Assoc. Comput. Linguistics, 7th Int. Joint Conf. Natural Language Process.*, vol. 1, 2015, pp. 167–176.
- [21] Y.-G. Cheong and R. M. Young, "A framework for summarizing game experiences as narratives," in *Proc. AIIDE*, 2006, pp. 106–108.
- [22] H. L. Chieu and Y. K. Lee, "Query based event extraction along a timeline," in *Proc. 27th Annu. Int. Conf. Res. Develop. Inf. Retr. (SIGIR)*, 2004, pp. 425–432.
- [23] G. L. Ciampaglia and J. Licato, "Fact-checking, false narratives, and argumentation schemes," in *Proc. CEUR Workshop*, 2021, pp. 1–18.
- [24] G. S. Collins, P. Dhiman, C. L. Andaur Navarro, J. Ma, L. Hooft, J. B. Reitsma, P. Logullo, A. L. Beam, L. Peng, B. Van Calster, M. van Smeden, R. D. Riley, and K. G. Moons, "Protocol for development of a reporting guideline (TRIPOD-AI) and risk of bias tool (PROBAST-AI) for diagnostic and prognostic prediction model studies based on artificial intelligence," *BMJ Open*, vol. 11, no. 7, Jul. 2021, Art. no. e048008.
- [25] S. R. Corman, "Understanding the role of narrative in extremist strategic communication," in *Countering Violent Extremism: Scientific Methods and Strategies*, vol. 36, Washington, DC, USA, 2011.
- [26] W. Croft, P. Peskova, and M. Regan, "Integrating decompositional event structures into storylines," in *Proc. Events Stories News Workshop*, 2017, pp. 98–109.
- [27] D. Das, D. Chen, A. F. T. Martins, N. Schneider, and N. A. Smith, "Frame-semantic parsing," *Comput. Linguistics*, vol. 40, no. 1, pp. 9–56, Mar. 2014.
- [28] S. Dathathri, A. Madotto, J. Lan, J. Hung, E. Frank, P. Molino, J. Yosinski, and R. Liu, "Plug and play language models: A simple approach to controlled text generation," 2019, *arXiv:1912.02164*.
- [29] S. Dey, "A persuasive approach to designing interactive tools around the promises and perils of social platforms," Ph.D. thesis, Dept. Comput. Sci., Univ. Illinois Urbana-Champaign, Champaign, IL, USA, 2020.
- [30] S. Dey, K. Karahalios, and W.-T. Fu, "Effects of socially stigmatized crowdfunding campaigns in shaping opinions," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Apr. 2018, pp. 1–13.
- [31] G. R. Doddington, A. Mitchell, M. A. Przybocki, L. A. Ramshaw, S. M. Strassel, and R. M. Weischedel, "The automatic content extraction (ACE) program-tasks, data, and evaluation," in *Proc. Lrec*, Lisbon, Portugal, vol. 2, 2004, pp. 837–840.
- [32] K. Egan, "What is a plot? *New Literary Hist.*, vol. 9, no. 3, pp. 455–473, 1978.
- [33] F. E. Outa, M. Francia, P. Marcel, V. Peralta, and P. Vassiliadis, "Supporting the generation of data narratives," in *Proc. ER Forum/Posters/Demos*, Nov. 2020, pp. 168–172.
- [34] K. Elkins and J. Chun, "Can GPT-3 pass a writer's Turing test?" *J. Cultural Anal.*, vol. 1, no. 1, p. 17212, 2020.
- [35] J. E. Escalas, "Imagine yourself in the product: Mental simulation, narrative transportation, and persuasion," *J. Advertising*, vol. 33, no. 2, pp. 37–48, Jan. 2004.
- [36] A. Fan, M. Lewis, and Y. Dauphin, "Hierarchical neural story generation," 2018, *arXiv:1805.04833*.
- [37] M. Fanton, H. Bonaldi, S. S. Tekiroglu, and M. Guerini, "Human-in-the-loop for data collection: A multi-target counter narrative dataset to fight online hate speech," 2021, *arXiv:2107.08720*.
- [38] D. Farmakiotou, V. Karkaletsis, J. Koutsias, G. Sigletos, C. D. Spyropoulos, and P. Stamatiopoulos, "Rule-based named entity recognition for Greek financial texts," in *Proc. Workshop Comput. Lexicography Multimedia Dictionaries (COMLEX)*, 2000, pp. 75–78.
- [39] F. Ferraro, T.-H. Huang, S. Lukin, and M. Mitchell, "Proceedings of the second workshop on storytelling," in *Proc. 2nd Workshop Storytelling*, 2019, pp. 1–10.
- [40] D. A. Fishman, "Review and criticism: Research pioneer tribute—Rethinking Marshall McLuhan: Reflections on a media theorist," *J. Broadcast. Electron. Media*, vol. 50, no. 3, pp. 567–574, Sep. 2006.
- [41] J. Frow, "The literary frame," *Narrative Dynamics: Essays on Time, Plot, Closure, and Frames*. Columbus, OH, USA: Ohio State Univ. Press, 2002, pp. 333–338.
- [42] G. P. C. Fung, J. X. Yu, H. Liu, and P. S. Yu, "Time-dependent event hierarchy construction," in *Proc. 13th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2007, pp. 300–309.
- [43] Y. Gabriel, "The Sage Handbook of Organizational Discourse, Chapter Narratives, Stories and Texts." New York, NY, USA: Sage, 2004, pp. 61–78.
- [44] L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy, "The pile: An 800GB dataset of diverse text for language modeling," 2021, *arXiv:2101.00027*.
- [45] A. Georgakopoulou, "Small stories research: A narrative paradigm for the analysis of social media," *The Sage Handbook of Social Media Research Methods*. New York, NY, USA: Sage, 2017, pp. 266–281.
- [46] P. Gervás, "Stories from games: Content and focalization selection in narrative composition," in *Proc. 1st Spanish Symp. Digit. Entertainment (SEED)*. Complutense Univ. of Madrid, Sep. 2013, p. 25.
- [47] S. Giannatos, M. Nelson, Y.-G. Cheong, and G. Yannakakis, "Suggesting new plot elements for an interactive story," in *Proc. AAAI Conf. Artif. Intell. Interact. Digit. Entertainment*, vol. 7, 2011, pp. 1230–1352.
- [48] J. Guan, F. Huang, Z. Zhao, X. Zhu, and M. Huang, "A knowledge-enhanced pretraining model for commonsense story generation," *Trans. Assoc. Comput. Linguistics*, vol. 8, pp. 93–108, Dec. 2020.
- [49] N. Guarino, D. Oberle, and S. Staab, "What is an ontology?" in *Handbook on Ontologies*. Cham, Switzerland: Springer, 2009, pp. 1–17.
- [50] J. F. Gubrium and J. A. Holstein, *Analyzing Narrative Reality*. New York, NY, USA: Sage, 2009.
- [51] T. Guha, N. Kumar, S. S. Narayanan, and S. L. Smith, "Computationally deconstructing movie narratives: An informatics approach," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Apr. 2015, pp. 2264–2268.
- [52] W. Guo and A. Caliskan, "Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.*, Jul. 2021, pp. 122–133.
- [53] A. H. Gutches, T. Hedden, S. Ketay, A. Aron, and J. D. E. Gabrieli, "Neural differences in the processing of semantic relationships across cultures," *Social Cognit. Affect. Neurosci.*, vol. 5, nos. 2–3, pp. 254–263, Jun. 2010.
- [54] N. Hassan, G. Zhang, F. Arslan, J. Caraballo, D. Jimenez, S. Gawsane, S. Hasan, M. Joseph, A. Kulkarni, A. K. Nayak, V. Sable, C. Li, and M. Tremayne, "ClaimBuster: The first-ever end-to-end fact-checking system," *Proc. VLDB Endowment*, vol. 10, no. 12, pp. 1945–1948, Aug. 2017.
- [55] E. Hovy, M. Marcus, M. Palmer, L. Ramshaw, and R. Weischedel, "OntoNotes: The 90% solution," in *Proc. 20th Hum. Lang. Technol. Conf. NAACL, Companion, Short Papers (NAACL)*, 2006, pp. 57–60.
- [56] J. Hu and P. Paroubek, "A fine-grained annotated corpus for target-based opinion analysis of economic and financial narratives," in *Proc. 3rd Workshop Econ. Natural Lang. Process.*, 2021, pp. 1–12.
- [57] Z. Hu, E. Rahimtooghi, and M. A. Walker, "Inference of fine-grained event causality from blogs and films," 2017, *arXiv:1708.09453*.
- [58] R. Huang, S. Ha, and S.-H. Kim, "Narrative persuasion in social media: An empirical study of luxury brand advertising," *J. Res. Interact. Marketing*, vol. 12, no. 3, pp. 274–292, Oct. 2018.
- [59] T.-H. Huang, F. Ferraro, N. Mostafazadeh, I. Misra, A. Agrawal, J. Devlin, R. Girshick, X. He, P. Kohli, and D. Batra, "Visual storytelling," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Lang. Technol.*, 2016, pp. 1233–1239.



- [60] M. N. Hussain, H. A. Rubaye, K. K. Bandeli, and N. Agarwal, "Stories from blogs: Computational extraction and visualization of narratives," in *Proc. Text2Story ECIR*, 2021, pp. 33–40.
- [61] M. N. Hussain, K. K. Bandeli, S. Al-khateeb, and N. Agarwal, "Analyzing shift in narratives regarding migrants in Europe via blogosphere," in *Proc. Text2Story ECIR*, 2018, pp. 33–40.
- [62] M. Iqbal, A. Karim, and F. Kamiran, "Bias-aware lexicon-based sentiment analysis," in *Proc. 30th Annu. ACM Symp. Appl. Comput.*, Apr. 2015, pp. 845–850.
- [63] H. Ji, P. Ke, S. Huang, F. Wei, X. Zhu, and M. Huang, "Language generation with multi-hop reasoning on commonsense knowledge graph," 2020, *arXiv:2009.11692*.
- [64] C. Jia, X. Liang, and Y. Zhang, "Cross-domain NER using Cross-domain language modeling," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, 2019, pp. 2464–2474.
- [65] X. Jin, F. Barbieri, B. Kennedy, A. M. Davani, L. Neves, and X. Ren, "On transferability of bias mitigation effects in language model fine-tuning," 2020, *arXiv:2010.12864*.
- [66] A. Jorge, R. Campos, A. Jatowt, and S. Nunes, "First international workshop on narrative extraction from texts," in *Proc. Text2Story Workshop ECIR*, 2018, pp. 833–834.
- [67] A. M. Jorge, R. Campos, A. Jatowt, and S. Bhatia, "The 2nd international workshop on narrative extraction from text: Text2story 2019," in *Proc. Eur. Conf. Inf. Retr. Cham, Switzerland: Springer*, 2019, pp. 389–393.
- [68] S. S. Kamath, A. Bagalkotkar, A. Kandelwal, S. Pandey, and K. Poornima, "Sentiment analysis based approaches for understanding user context in web content," in *Proc. Int. Conf. Commun. Syst. Netw. Technol.*, Apr. 2013, pp. 607–611.
- [69] S. Mellace and A. Antonucci, "Temporal embeddings and transformer models for narrative text understanding," 2020, *arXiv:2003.08811*.
- [70] K. T. Kernan, "Semantic relationships and the child's acquisition of language," *Anthropol. Linguistics*, vol. 12, no. 5, pp. 171–187, May 1970.
- [71] N. Shirish Kesar, B. McCann, L. R. Varshney, C. Xiong, and R. Socher, "CTRL: A conditional transformer language model for controllable generation," 2019, *arXiv:1909.05858*.
- [72] C. S. G. Khoo and J. C. Na, "Semantic relations in information science," *Inf. Sci. Technol.*, vol. 40, no. 1, pp. 157–228, 2006.
- [73] N. M. Krause, I. Freiling, B. Beets, and D. Brossard, "Fact-checking as risk communication: The multi-layered risk of misinformation in times of COVID-19," *J. Risk Res.*, vol. 23, nos. 7–8, pp. 1052–1059, Aug. 2020.
- [74] A. Kutuzov, E. Velldal, and L. Øvrelid, "Tracing armed conflicts with diachronic word embedding models," in *Proc. Events Stories News Workshop*, 2017, pp. 31–36.
- [75] V. Labatut and X. Bost, "Extraction and analysis of fictional character networks: A survey," *ACM Comput. Surv.*, vol. 52, no. 5, pp. 1–40, Sep. 2020.
- [76] G. Lakoff and S. Narayanan, "Toward a computational model of narrative," in *Proc. AAAI Fall Symp.*, 2010, pp. 1–19.
- [77] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, and C. Dyer, "Neural architectures for named entity recognition," 2016, *arXiv:1603.01360*.
- [78] A. Langer, M.-A. Kaufhold, E. M. Runft, C. Reuter, M. Grinko, and V. Pipek, "Counter narratives in social media: An empirical study on combat and prevention of terrorism," in *Proc. ISCRAM*, 2019, pp. 1–10.
- [79] J. Lankow, J. Ritchie, and R. Crooks, *Infographics: The Power of Visual Storytelling*. Hoboken, NJ, USA: Wiley, 2012.
- [80] J.-S. Lee and J. Hsiang, "Patent claim generation by fine-tuning openai GPT-2," *World Patent Inf.*, vol. 62, May 2020, Art. no. 101983.
- [81] O.-J. Lee and J. J. Jung, "Story embedding: Learning distributed representations of stories based on character networks," *Artif. Intell.*, vol. 281, Apr. 2020, Art. no. 103235.
- [82] J. Leng and P. Jiang, "A deep learning approach for relationship extraction from interaction context in social manufacturing paradigm," *Knowl.-Based Syst.*, vol. 100, pp. 188–199, May 2016.
- [83] E. Levi, G. Mor, S. Shenhav, and T. Sheaffer, "CompRes: A dataset for narrative structure in news," 2020, *arXiv:2007.04874*.
- [84] B. Li, H. Peng, R. Sainju, J. Yang, L. Yang, Y. Liang, W. Jiang, B. Wang, H. Liu, and C. Ding, "Detecting gender bias in transformer-based models: A case study on BERT," 2021, *arXiv:2110.15733*.
- [85] J. Li, A. Sun, J. Han, and C. Li, "A survey on deep learning for named entity recognition," *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 1, pp. 50–70, Jan. 2022.
- [86] S. Li, "Using language models in causal story generation," M.S. thesis, Dept. Elect. Eng., Georgia Inst. Technol., Atlanta, GA, USA, 2021.
- [87] F.-R. Lin and C.-H. Liang, "Storyline-based summarization for news topic retrospection," *Decis. Support Syst.*, vol. 45, no. 3, pp. 473–490, 2008.
- [88] S. Liu, Y. Chen, S. He, K. Liu, and J. Zhao, "Leveraging framenet to improve automatic event detection," in *Proc. 54th Annu. Meeting Assoc. Comput. Linguistics*, vol. 1, 2016, pp. 2134–2143.
- [89] Y. Liu, Z. Shi, and A. Sarkar, "Exploiting rich syntactic information for relation extraction from biomedical articles," in *Proc. Hum. Lang. Technol., Conf. North Amer. Chapter Assoc. Comput. Linguistics, Companion Volume, Short Papers (NAACL)*, 2007, pp. 97–100.
- [90] R. K. Logan, "The five ages of communication," *Explor. Media Ecol.*, vol. 1, no. 1, pp. 13–20, Apr. 2002.
- [91] L. Lucy and D. Bamman, "Gender and representation bias in GPT-3 generated stories," in *Proc. 3rd Workshop Narrative Understand.*, 2021, pp. 48–55.
- [92] M. A. Madaio, L. Stark, J. Wortman Vaughan, and H. Wallach, "Co-designing checklists to understand organizational challenges and opportunities around fairness in AI," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Apr. 2020, pp. 1–14.
- [93] L. Magee, L. Ghahremanlou, K. Soldatic, and S. Robertson, "Intersectional bias in causal language models," 2021, *arXiv:2107.07691*.
- [94] H. Henry Mao, B. P. Majumder, J. McAuley, and G. W. Cottrell, "Improving neural story generation by targeted common sense grounding," 2019, *arXiv:1908.09451*.
- [95] T. Marcoux and N. Agarwal, "Narrative trends of COVID-19 misinformation," in *Proc. Text2Story ECIR*, 2021, pp. 77–80.
- [96] L. Martin, P. Ammanabrolu, X. Wang, W. Hancock, S. Singh, B. Harrison, and M. Riedl, "Event representations for automated story generation with deep neural nets," in *Proc. AAAI Conf. Artif. Intell.*, vol. 32, 2018, pp. 189–196.
- [97] C. May, A. Wang, S. Bordia, S. R. Bowman, and R. Rudinger, "On measuring social biases in sentence encoders," 2019, *arXiv:1903.10561*.
- [98] R. McCreadie, R. L. T. Santos, C. Macdonald, and I. Ounis, "Explicit diversification of event aspects for temporal summarization," *ACM Trans. Inf. Syst.*, vol. 36, no. 3, pp. 1–31, Apr. 2018.
- [99] Q. Mei and C. Zhai, "Discovering evolutionary theme patterns from text: An exploration of temporal text mining," in *Proc. 11th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2005, pp. 198–207.
- [100] Y. Mori, H. Yamane, Y. Mukuta, and T. Harada, "The nectar of missing position prediction for story completion," in *Proc. Text2Story ECIR*, 2021, pp. 65–69.
- [101] N. Mostafazadeh, N. Chambers, X. He, D. Parikh, D. Batra, L. Vanderwende, P. Kohli, and J. Allen, "A corpus and evaluation framework for deeper understanding of commonsense stories," 2016, *arXiv:1604.01696*.
- [102] N. Mostafazadeh, M. Roth, A. Louis, N. Chambers, and J. Allen, "LSD-Sem 2017 shared task: The story cloze test," in *Proc. 2nd Workshop Linking Models Lexical, Sentential Discourse-Level Semantics*, 2017, pp. 46–51.
- [103] D. Nadeau and S. Sekine, "A survey of named entity recognition and classification," *Linguisticae Investigationes*, vol. 30, no. 1, pp. 3–26, Jan. 2007.
- [104] R. Nallapati, A. Feng, F. Peng, and J. Allan, "Event threading within news topics," in *Proc. 13th ACM Conf. Inf. Knowl. Manag. (CIKM)*, 2004, pp. 446–453.
- [105] R. Nallapati, B. Zhou, C. dos Santos, C. Gulcehre, and B. Xiang, "Abstractive text summarization using sequence-to-sequence RNNs and beyond," in *Proc. 20th SIGNLL Conf. Comput. Natural Lang. Learn.*, 2016, pp. 280–290.
- [106] G. Nathaniel, "Solving the story cloze test using graph feature extraction," in *Intelligent Computing*. Cham, Switzerland: Springer, 2022, pp. 1000–1013.
- [107] C. D. Nguyen, E. Gallagher, A. Read, and G.-J. D. Vreede, "Generating user stories in groups," in *Proc. Int. Conf. Collaboration Technol.* Cham, Switzerland: Springer, 2009, pp. 357–364.
- [108] B. Osatuyi, "Information sharing on social media sites," *Comput. Hum. Behav.*, vol. 29, no. 6, pp. 2622–2631, Nov. 2013.
- [109] P. Oza and L. Dietz, "Which entities are relevant for the story?" in *Proc. Text2Story, Workshop Narrative Extraction Texts, Eur. Conf. Inf. Retr.*, 2021, pp. 41–48.
- [110] K. Padian, "Narrative and 'anti-narrative' in science: How scientists tell stories, and don't," *Integrative Comparative Biol.*, vol. 58, no. 6, pp. 1224–1234, 2018.

- [111] R. Page, *The Handbook of Narrative Analysis, Chapter The 17 Narrative Dimensions of Social Media Storytelling*. Hoboken, NJ, USA: Wiley, 2015, pp. 329–347.
- [112] R. Page, *Narratives Online: Shared Stories in Social Media*. Cambridge, U.K.: Cambridge Univ. Press, 2018.
- [113] R. E. Page, *Stories and Social Media: Identities and Interaction*. London, U.K.: Routledge, 2013.
- [114] C. Painter, J. Martin, and L. Unsworth, *Reading Visual Narratives: Image Analysis of Children's Picture Books*. Sheffield, U.K.: Equinox, 2013.
- [115] A. Pasquali, V. Mangaravite, R. Campos, A. M. Jorge, and A. Jatowt, "Interactive system for automatically generating temporal narratives," in *Proc. Eur. Conf. Inf. Retr. Cham, Switzerland: Springer*, 2019, pp. 251–255.
- [116] A. Pingle, A. Piplai, S. Mittal, A. Joshi, J. Holt, and R. Zak, "RelExt: Relation extraction using deep learning approaches for cybersecurity knowledge graph improvement," in *Proc. IEEE/ACM Int. Conf. Adv. Social Netw. Anal. Mining*, Aug. 2019, pp. 879–886.
- [117] A. Piplai, P. Ranade, A. Kotal, S. Mittal, S. N. Narayanan, and A. Joshi, "Using knowledge graphs and reinforcement learning for malware analysis," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2020, pp. 2626–2633.
- [118] S. Powers and M. Kounalakis, "Can public diplomacy survive the internet? Bots, echo chambers, and disinformation," *Digit. Commons*, vol. 377, Mar. 2017. [Online]. Available: <https://digitalcommons.unl.edu/usdoepub/377/>
- [119] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Improving language understanding by generative pre-training," 2018. [Online]. Available: [https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language\\_understanding\\_paper.pdf](https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf)
- [120] P. Ranade, A. Piplai, S. Mittal, A. Joshi, and T. Finin, "Generating fake cyber threat intelligence using transformer-based models," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, 2021, pp. 12–18.
- [121] L. Ratinov and D. Roth, "Design challenges and misconceptions in named entity recognition," in *Proc. 13th Conf. Comput. Natural Lang. Learn. (CoNLL)*, 2009, pp. 147–155.
- [122] G. Rehm, J. M. Schneider, P. Bourgonje, A. Srivastava, J. Nehring, A. Berger, L. König, S. Räuchle, and J. Gerth, "Event detection and semantic storytelling: Generating a travelogue from a large collection of personal letters," in *Proc. Events Stories News Workshop*, 2017, pp. 42–51.
- [123] M. T. Ribeiro, T. Wu, C. Guestrin, and S. Singh, "Beyond accuracy: Behavioral testing of NLP models with checklist," 2020, *arXiv:2005.04118*.
- [124] A. Roberts, R. Gaizauskas, and M. Hepple, "Extracting clinical relationships from patient narratives," in *Proc. Workshop Current Trends Biomed. Natural Lang. Process. (BioNLP)*, 2008, pp. 10–18.
- [125] S. Sandberg, "The importance of stories untold: Life-story, event-story and trope," *Crime, Media, Culture, Int. J.*, vol. 12, no. 2, pp. 153–171, Aug. 2016.
- [126] E. F. T. K. Sang and F. De Meulder, "Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition," 2003, *arXiv:cs/0306050*.
- [127] S. Sarawagi, *Information Extraction*. Norwell, MA, USA: Now, 2008.
- [128] E. Segel and J. Heer, "Narrative visualization: Telling stories with data," *IEEE Trans. Vis. Comput. Graphics*, vol. 16, no. 6, pp. 1139–1148, Nov. 2010.
- [129] R. Segers, T. Caselli, and P. Vossen, "The circumstantial event ontology (CEO)," in *Proc. Events Stories News Workshop*, 2017, pp. 37–41.
- [130] D. Shahaf and C. Guestrin, "Connecting the dots between news articles," in *Proc. 16th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2010, pp. 623–632.
- [131] M. Shang, Z. Fu, H. Yin, B. Tang, D. Zhao, and R. Yan, "Find a reasonable ending for stories: Does logic relation help the story cloze test?" in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 10031–10032.
- [132] C. Shao, G. L. Ciampaglia, A. Flammini, and F. Menczer, "Hoaxy: A platform for tracking online misinformation," in *Proc. 25th Int. Conf. Companion World Wide Web (WWW Companion)*, 2016, pp. 745–750.
- [133] R. Sharma, J. Allen, O. Bakhshandeh, and N. Mostafazadeh, "Tackling the story ending biases in the story cloze test," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 752–757.
- [134] S. Shekarpour, F. Alshargi, K. Thirunaravan, V. L. Shalin, and A. Sheth, "CEVO: Comprehensive EVent ontology enhancing cognitive annotation on relations," in *Proc. IEEE 13th Int. Conf. Semantic Comput. (ICSC)*, Jan. 2019, pp. 385–391.
- [135] S. Shekarpour, A. Saxena, K. Thirunarayan, V. L. Shalin, and A. Sheth, "Principles for developing a knowledge graph of interlinked events from news headlines on Twitter," 2018, *arXiv:1808.02022*.
- [136] A. Sheth, I. B. Arpinar, and V. Kashyap, "Relationships at the heart of semantic web: Modeling, discovering, and exploiting complex semantic relationships," in *Enhancing the Power of the Internet*. Cham, Switzerland: Springer, 2004, pp. 63–94.
- [137] B. Singh and H. K. Singh, "Web data mining research: A survey," in *Proc. IEEE Int. Conf. Comput. Intell. Comput. Res.*, Dec. 2010, pp. 1–10.
- [138] E. Spiliopoulou, E. Hovy, and T. Mitamura, "Event detection using frame-semantic parser," in *Proc. Events Stories News Workshop*, 2017, pp. 15–20.
- [139] T. Spinde, J. Krieger, T. Ruas, J. Mitrovic, F. Götz-Hahn, A. Aizawa, and B. Gipp, "Exploiting transformer-based multitask learning for the detection of media bias in news articles," in *Proc. Int. Conf. Inf. Cham, Switzerland: Springer*, Feb. 2022, pp. 225–235.
- [140] P. Stenetorp, S. Pyysalo, G. Topić, T. Ohta, S. Ananiadou, and J. Tsujii, "BRAT: A web-based tool for nlp-assisted text annotation," in *Proc. Demonstrations, 13th Conf. Eur. Chapter Assoc. Comput. Linguistics*, 2012, pp. 102–107.
- [141] V. C. Storey, "Understanding semantic relationships," *Vldb J.*, vol. 2, no. 4, pp. 455–488, Oct. 1993.
- [142] W. L. Taylor, "'Cloze procedure': A new tool for measuring readability," *Journalism Quart.*, vol. 30, no. 4, pp. 415–433, 1953.
- [143] M. Thelwall, "Gender bias in sentiment analysis," *Online Inf. Rev.*, vol. 42, no. 1, pp. 45–57, Feb. 2018.
- [144] A. Tikhonov, I. Samenko, and I. P. Yamshchikov, "StoryDB: Broad multi-language narrative dataset," 2021, *arXiv:2109.14396*.
- [145] J. Li, C. Zhu, S. Li, T. Zhao, and D. Zheng, "Exploiting Wikipedia priori knowledge for Chinese named entity recognition," in *Proc. 12th Int. Conf. Natural Comput., Fuzzy Syst. Knowl. Discovery (ICNC-FSKD)*, Aug. 2016, pp. 698–707.
- [146] M. M. Tuffield, D. E. Millard, and N. R. Shadbolt, "Ontological approaches to modelling narrative," in *Proc. 2nd AKT DTA Symp.*, 2006, pp. 45–54.
- [147] K. V. Dalen-Oskam, J. D. Does, M. Marx, I. Sijaranamual, K. Depuydt, B. Verheij, and V. Geirnaert, "Named entity recognition and resolution for literary studies," *Comput. Linguistics Netherlands J.*, vol. 4, pp. 121–136, May 2014.
- [148] P. Wang, N. Peng, F. Ilievski, P. Szekely, and X. Ren, "Connecting the dots: A knowledgeable path generator for commonsense question answering," 2020, *arXiv:2005.00691*.
- [149] A. Whiting and D. Williams, "Why people use social media: A uses and gratifications approach," *Qualitative Market Res., Int. J.*, vol. 16, no. 4, pp. 362–369, Aug. 2013.
- [150] M. E. Winston, R. Chaffin, and D. Herrmann, "A taxonomy of part-whole relations," *Cognit. Sci.*, vol. 11, no. 4, pp. 417–444, Oct. 1987.
- [151] W. Xiang and B. Wang, "A survey of event extraction from text," *IEEE Access*, vol. 7, pp. 173111–173137, 2019.
- [152] Y. Xie, Y. Hu, L. Xing, C. Wang, Y. Hu, X. Wei, and Y. Sun, "Enhancing pre-trained language models by self-supervised learning for story cloze test," in *Proc. Int. Conf. Knowl. Sci., Eng. Manag. Cham, Switzerland: Springer*, 2020, pp. 271–279.
- [153] R. Yan, X. Wan, J. Otterbacher, L. Kong, X. Li, and Y. Zhang, "Evolutionary timeline summarization: A balanced optimization framework via iterative substitution," in *Proc. 34th Int. ACM SIGIR Conf. Res. Develop. Inf. (SIGIR)*, 2011, pp. 745–754.
- [154] X. Yan, A. Naik, Y. Jo, and C. Rose, "Using functional schemas to understand social media narratives," in *Proc. 2nd Workshop Storytelling*, 2019, pp. 22–33.
- [155] J. Yin, X. Jiang, Z. Lu, L. Shang, H. Li, and X. Li, "Neural generative question answering," in *Proc. Workshop Hum.-Comput. Question Answering*, 2016, pp. 2972–2978.
- [156] W. Zaghouani, "RENAR: A rule-based Arabic named entity recognition system," *ACM Trans. Asian Lang. Inf. Process.*, vol. 11, no. 1, pp. 1–13, Mar. 2012.
- [157] R. Zellers, A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi, "Defending against neural fake news," 2019, *arXiv:1905.12616*.
- [158] Y. Zhang, S. Sun, M. Galley, Y.-C. Chen, C. Brockett, X. Gao, J. Gao, J. Liu, and B. Dolan, "DialoGPT: Large-scale generative pre-training for conversational response generation," 2019, *arXiv:1911.00536*.
- [159] X. Zhu and T. Oates, "Finding story chains in newswire articles," in *Proc. IEEE 13th Int. Conf. Inf. Reuse Integr. (IRI)*, Aug. 2012, pp. 93–100.




**PRIYANKA RANADE** is currently pursuing the Ph.D. degree in computer science with the University of Maryland, Baltimore County. She is a Research Scientist at the Laboratory for Physical Sciences and an Adjunct Professor at the ACES Program, University of Maryland, College Park. Her primary research interests include cybersecurity, knowledge representation, and social media analytics.



**SANORITA DEY** is currently an Assistant Professor with the Computer Science and Electrical Engineering Department, University of Maryland, Baltimore County (UMBC), where she studies human-computer interaction (HCI). Her research interests include social computing, crowd computing, technologically mediated persuasive systems, and spatial learning. Her work has been published and received awards in premier venues in human-computer interaction, including CHI and CSCW. She was a recipient of the Maria Pia Graton Fellowship, the UIUC Graduate College Dissertation Completion Fellowship, and the Outstanding Thesis Awards for her master's thesis. Her recent research on "Effects of Socially Stigmatized Crowdfunding Campaigns in Shaping Opinions" received the Best Paper Honorable Mention Award at CHI'18. Her other work on "Understanding the Effect of the Combination of Navigation Tools in Learning Spatial Knowledge" was selected among the Top-5 papers at SUI'19.



**ANUPAM JOSHI** (Fellow, IEEE) received the B.Tech. degree from IIT Delhi, in 1989, and the master's and Ph.D. degrees from Purdue University, in 1991 and 1993, respectively. He is currently the Oros Family Professor and the Chair of the Computer Science and Electrical Engineering Department, University of Maryland, Baltimore County (UMBC). He is also the Director of UMBC's Center for Cybersecurity. He has published over 275 technical papers with an H-index of 87 and over 28500 citations (per Google scholar), has been granted nine patents, and has obtained research support from the National Science Foundation (NSF), NASA, the Defense Advanced Research Projects Agency (DARPA), U.S. Department of Defense (DoD), NIST, IBM, Microsoft, Qualcomm, Northrop Grumman, and Lockheed Martin amongst others. His research interests include networked computing and intelligent systems. His primary focus has been on data management and security/privacy in mobile/pervasive computing environments, and policy driven approaches to security and privacy. He is also interested in semantic web and data/text/web analytics, especially their applications to (cyber) security and healthcare.



**TIM FININ** received the degrees from MIT and the University of Illinois at Urbana-Champaign. He has held positions at UMBC, Unisys, the University of Pennsylvania, Johns Hopkins University, and the MIT AI Laboratory. He is currently the Willard and Lillian Hackerman Chair in engineering and a Professor of computer science and electrical engineering at UMBC. He has over 40 years of experience in applying artificial intelligence to information systems and language understanding problems. His current research interests include representing and reasoning with knowledge graphs, extracting information from text, and cybersecurity. He is a fellow of ACM and AAAI and an IEEE Technical Achievement Award Recipient.

...