

Lost in Recursion: Mining Rich Event Semantics in Knowledge Graphs

Florian Plötzky
ploetzky@ifis.cs.tu-bs.de
Institute for Information Systems
TU Braunschweig
Braunschweig, Germany

Niklas Kiehne
kiehne@ifis.cs.tu-bs.de
Institute for Information Systems
TU Braunschweig
Braunschweig, Germany

Wolf-Tilo Balke
balke@ifis.cs.tu-bs.de
Institute for Information Systems
TU Braunschweig
Braunschweig, Germany

ABSTRACT

Our world is shaped by events of various complexity. This includes both small-scale local events like local farmer markets and large complex events like political and military conflicts. The latter are typically not observed directly but through the lenses of intermediaries like newspapers or social media. In other words, we do not witness the unfolding of such events directly but are confronted with *narratives* surrounding them. Such narratives capture different aspects of a complex event and may also differ with respect to the *narrator*. Thus, they provide a rich semantics concerning real-world events. In this paper, we show how narratives concerning complex events can be constructed and utilized. We provide a formal representation of narratives based on *recursive nodes* to represent multiple levels of detail and discuss how narratives can be bound to event-centric knowledge graphs. Additionally, we provide an algorithm based on incremental prompting techniques that mines such narratives from texts to account for different perspectives on complex events. Finally, we show the effectiveness and future research directions in a proof of concept.

CCS CONCEPTS

• Information systems → Web mining; • Computing methodologies → Natural language processing.

KEYWORDS

Narratives, Events, Recursive Narrative Mining

ACM Reference Format:

Florian Plötzky, Niklas Kiehne, and Wolf-Tilo Balke. 2024. Lost in Recursion: Mining Rich Event Semantics in Knowledge Graphs. In *ACM Web Science Conference (WebSci '24)*, May 21–24, 2024, Stuttgart, Germany. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3614419.3644001>

1 INTRODUCTION

With the rise of the Web, research concerning large and complex events has become necessary to connect the dots and gain insights into the intricacies of such events [20]. However, most events that occur are observed through the lenses of Web content like online news and social media. As an effect, we mostly do not observe the

event itself but are presented with *narratives* about it. Consequently, journalists, analysts, and political scientists are often confronted with different and potentially incompatible stories regarding a specific event, which must be compared and unified [2, 10]. Such tasks are aggravated by the usage of framing techniques [10, 27] and patterns of strategic communication [47]. In either case, systems enabling the representation and comparison of narratives regarding complex events are beneficial to uncover different perspectives on such events and unpack the surrounding narratives.

At the core, we understand a narrative as a sequence of events arranged by a *narrator* to describe properties and inner workings of complex events. For instance, an event like a war usually comprises different battles, an event that started the war (which is not necessarily a battle), and an event that ended it. Additionally, narratives can be used to describe properties and justify specific events. Take, for instance, the narrative told by Vladimir Putin in his address to the nation at the advent of the Ukraine invasion in 2022.¹ Here, Putin connects over 30 years of events in a specific way, introduces new events like the redivision of the world (i.e., different wars involving Western countries since the Collapse of the Soviet Union), and frames the role of participants like NATO in a specific way. Fig. 1 shows an excerpt of this narrative in a structured way that will also be our running example throughout the paper.

Ultimately, this whole narrative is crafted to justify the invasion. Obviously, parts of it are framed and adjusted to achieve this objective (cf., [9] for a critical discussion and a different point of view). Hence, a narrative constructed from a more Western perspective like that of the United States would be different at large, especially regarding the motivation of the NATO eastward expansion in that narrative. Still, both narratives would concern the same named event, i.e., the Russo-Ukrainian War, regarding the war's cause. The differences and similarities of such narratives, however, may expose specific highlighting of events and, therefore, hint at the intentions of the respective narrator.

In this paper, we focus on how narratives of complex events can be used to describe the inner structure of an event from different points of view (i.e., by different narrators). Specifically, we refer to the idea of the faculty of language [14] to construct narrative events chains of different granularity. The core idea is that since human language is constructed recursively, narratives can also be seen as the result of compositions of different narratives, where a narrative is simply the composition of at least two events [18]. Putin's narrative is a composition of narratives concerning the collapse of the Soviet Union, the events around the NATO eastward

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
WebSci '24, May 21–24, 2024, Stuttgart, Germany
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0334-8/24/05.
<https://doi.org/10.1145/3614419.3644001>

¹<https://www.bloomberg.com/news/articles/2022-02-24/full-transcript-vladimir-putin-s-televized-address-to-russia-on-ukraine-feb-24>, last accessed Feb. 15, 2024




Figure 1: A simplified narrative that depicts Vladimir Putin’s address to the nation. The events “Redivision of the World” and “Iraq War” are recursive nodes in this example and reflect the inner structure of those events as exemplified in Putin’s speech. All events also available in Wikidata are marked with their respective identifier.

expansion, and the redivision of the world as depicted in Fig. 1. All of those different pieces must, of course, fit (from the narrator’s view) to create a unified narrative of the particular event in question, i.e., the Russo-Ukrainian War.

Of course, representing complex events is not a new topic. Most research in this direction was conducted in works concerning ontological aspects of events [1, 13, 41] and put in practice by the implementation of event-centric knowledge graphs (ECKGs) like EventKG [11] and OEKG [12]. While some recent works introduced narrative aspects [31, 33] and viewpoint-specific annotations [32], narratives are not only characterized by role-assignments to participants [10] but also by the choice and arrangement the events discussed in it. Putin’s speech, for instance, arranges a sequence of events as an explanatory factor for the invasion of Ukraine (from his perspective). However, a corresponding ECKG may not directly reflect the narrative. That is, a narrative may combine otherwise unrelated events, for instance in a causal relationship, with respect to a specific *viewpoint*. Therefore, such relationships are only valid from a specific point of view but not generally agreed upon, and hence, they are not materialized in ECKGs.

In addition, narratives can be organized hierarchically using *recursive nodes*, e.g., “Redivision of the World” involves three distinct and usually unconnected events. Another application for such recursive nodes is to capture a specific aspect a narrator wants to highlight. For our running example, this is true for the Iraq War, where the (in hindsight false) allegations of the Iraqi possession of weapons of mass destruction (WMD) in a meeting of the UN Security Council are explicitly highlighted [3]. A narrative can be seen as an *overlay* for ECKGs that connects and constructs subgraphs from the KG and enriches it with additional (potentially missing) events only available in the narrative. By grounding narratives in

well-known ECKGs, we can express and compare different perspectives on complex events and reason about, for instance, similarities, differences, and variants of different narratives. Hence, we argue that narratives add a rich semantic to ECKGs.

In this paper, we adopt and extend the narrative model of Kroll et al. [18] specifically to represent narratives concerning real-world events in a recursive manner. We contribute a method to mine narratives from different perspectives and on multiple levels based on state-of-the-art event schema induction techniques [25]. Additionally, we discuss how such mined events should be linked to ECKGs. Finally, we demonstrate the effectiveness of our approach in a proof of concept.

2 EVENTS AND NARRATIVES

In this section, we first introduce events, ECKGs, and viewpoints to represent narrative aspects of ECKGs. We then provide a formalization of narratives, which is used throughout this paper.

2.1 Events

In general, events connect entities (the participants of the event) and set them into a temporal and spatial context. We define events in a formal way as follows:

DEFINITION 1 (EVENTS). *Given a set of entities $\mathcal{E}$, a set of locations $\mathcal{L}$ (which could either be entities or literal values), and literals referring to points in time $\mathcal{T}$, an event is an interaction between participants $p \in \mathcal{E}$ that takes place at a given location $l \in \mathcal{L}$ and at a specific point in time $t \in \mathcal{T}$ or in a specific time span $(t_1, t_2) \in \mathcal{T} \times \mathcal{T}$. The set of events is denoted by Γ .*

Optionally, events can be described by an *event type* that describes the typical shape of an event, e.g., what kind of participants are involved.

Of course, this definition is quite broad. That is, it can fit a small-scale event like a local farmer’s market with a few hours duration and a year-long war like the ongoing war between Russia and Ukraine. The latter events are often drilled down to multiple, simpler ones. For the Russo-Ukrainian War we can, for instance, consider the Snake Island Campaign (Q111012552) or the Battle of Chernihiv (Q111013915), both in 2022, as parts of this event. To account for such complex events, we define them accordingly.

DEFINITION 2 (COMPLEX EVENTS). *A complex event is an event that is connected to any other event by either a mereological relationship (i.e., a part-of relation) or a hierarchical relationship (i.e., a sub-event relation).*

Current ECKGs primarily reflect on complex events by sub-event relationships as first introduced on a larger scale in EventKG [11]. For each event that is part of a complex event, the participants, location, and time must fit into the analogous components of the complex event. That is if the complex event has a time span of one year, the time of each sub-event must be located in this span. With those definitions in place we can formally introduce event-centric knowledge graphs.

DEFINITION 3 (EVENT-CENTRIC KNOWLEDGE GRAPHS). *An Event-Centric Knowledge Graph is a knowledge graph (i.e., a triplestore consisting of subject-predicate-object triples) where the majority of triples contain a (complex) event in their subject or object respectively.*

Typically, such knowledge graphs are constructed by using the well-known Resource Description Framework [17]. Besides knowledge graphs explicitly designed as ECKGs (e.g., EventKG and OEKG), this definition also encompasses portions of general-purpose knowledge graphs like Wikidata [46]. The latter case was already illustrated in [40]. Events in such graphs can be seen as a special form of an entity in the graph that adheres to Def. 1, i.e., that connects other entities and bounds this interaction in space and time. Entities are included in such graphs mostly as participants for events. Additionally, they can be further characterized by additional attributes that are not bound to their participation in an event (e.g., typical attributes like names, birth dates, and nationalities).

2.2 Viewpoints

ECKGs offer structured event representation at scale. In contrast to textual event repositories (e.g., GDELT [21] and Event Registry [20]), they also offer a canonical representation of named events and allow for querying and reasoning about events. However, transforming event knowledge from unstructured sources into ECKGs tends to miss *narrative aspects* for events, e.g., specific roles like an *aggressor* in a conflict. Such aspects can be problematic since it might be disputed whether an event participant can be seen as, for instance, an aggressor. Tackling this problem, we recently introduced *viewpoints*, i.e., the representation of disputed facts from multiple perspectives to ECKGs [32]. We rely on this work and define viewpoints as follows.

DEFINITION 4 (VIEWPOINTS). *A viewpoint represents the consensual stance (valid or invalid) towards a specific event-centric claim that is expressed by a group of entities. Each entity in such a group of entities G must be capable of expressing a stance. We denote the set of viewpoints as $\mathcal{V}$.*

At the core, specific claims concerning an event (e.g., the role a participant played in the event) might be disputed, in terms of their validity, by different groups of entities G . A viewpoint aggregates the individual stances of G regarding the validity of this claim into a consensual stance for this group, e.g., by a majority vote. Examples for typical group members include politicians, newspaper columnists, and speakers of governmental and non-governmental organizations.

Viewpoints belong to a specific viewpoint model that determines whether two viewpoints $v_1, v_2 \in \mathcal{V}$ can be subsumed, i.e., whether a claim that is valid in v_1 can also be seen as valid in v_2 . A simple variant of such a viewpoint model is a group hierarchy that determines which groups G_1, G_2 can be aggregated to form a larger group G^* . For instance, the stances of the President of the United States and the US Congress may be aggregated into the “US viewpoint”. If two viewpoints v_1 and v_2 can be aggregated according to such a model, we call them *viewpoint-compatible*. Additionally, we call predicates in ECKGs that are annotated with viewpoints *attributions*. An example is depicted in Fig. 2 between Event 3 and an entity.

2.3 Narratives

In this section, we finally introduce narratives as a representation to connect events, for instance, in ECKGs meaningfully. As a motivation, we refer to the recent idea of hierarchical event schemas as presented in [25]. Basically, complex event types like wars or pandemic outbreaks are described by hierarchical schema graphs. A schema graph connects the most important entities and sub-events of an event type to give an abstract description of what typically happens in those scenarios. This can be seen as a hierarchical schema since each sub-event also spans an event schema graph that can be traversed recursively downwards.

We argue that recursion in such schemas is a valuable mechanism for the deep representation of otherwise complex events. Especially in conflict situations, understanding the root causes of how the actual conflict played out can help researchers understand the actual mechanism behind the conflict for purposes of categorization and comparison with other events.

In comparison, narratives enhance this idea in two aspects. First, while event schemas aim to describe the schema of a typical representative of an event, real-world events are often convoluted or surrogate events. Take, for instance, the “Redivision of the World” in Fig. 1. This event comprises a number of real-world events but is constructed in the narrative itself. Hence, there will be no counterpart in an ECKG besides the sub-events like the Belgrade invasion, which can be found in ECKGs and other event-centric repositories. In addition, convoluted events are those that do not follow a typical schema. This can be seen by the NATO eastward expansion that comprises multiple steps and can hardly be categorized by a general schema. Second and in addition to Sect. 2.2, especially large-scale events like wars, the choice and connotation of events depends on the narrator up to a certain degree. That is, depending on the perspective chosen, specific events are highlighted or framed in a way. In Fig. 1, this is visible for the Iraq War. While this year-long war comprises an extensive range of events, the narrator chose to




Figure 2: Narratives and their bindings to ECKGs.

present the WMD misinformation of Colin Powell as the precursor of the invasion. One can argue why this is the case, but we argue that different narratives on the same event yield insights into consensual knowledge regarding the event.

In this paper we use and enhance [18] as a basis for our narrative definition.

DEFINITION 5 (NARRATIVES). A narrative is defined as follows:

- (1) Let L be a set of literals. We call a directed edge-labeled graph (V, E) a narrative with $V \subseteq \mathcal{E} \cup \Gamma \cup L$ being the nodes and $E \subseteq \mathcal{R}_F \cup \mathcal{R}_N$ being the edges. $\mathcal{R}_F$ is the set of factual relationships that connect events and entities with entities and literals. $\mathcal{R}_N$ is the set of narrative relationships that connect events with events. We denote the set of narratives as $\mathcal{N}$.
- (2) For all $n \in \mathcal{N}$ we define a function $\eta(e, n)$ with $\eta : \Gamma \times \mathcal{N} \rightarrow \mathcal{N}$ that returns the narrative for event e in n .

For our purpose, we can model narratives as directed graphs. The nodes of the graphs are entities denoted by $\mathcal{E}$, event labels represented by Γ , and literals denoted by L . Hence, narratives are able to connect events and their participants and enrich them with additional properties (expressed by labels). We define two types of relationships. Factual relationships, denoted by $\mathcal{R}_F$, are used to connect events and entities to entities or literals. They are used to describe relationships and roles of participants to events, and to enrich them by using literals. In addition, we define relationships between events and events as narrative relationships, denoted by $\mathcal{R}_N$. Such relationships describe the *unfolding of events* in the narrative, e.g., temporal and causal relationships.

Each event $e \in \Gamma$ in a narrative $n \in \mathcal{N}$ is represented by a unique *event label* with a concise description of the event. Note that this label does not necessarily correspond to an event in an ECKG. Referring to Fig. 1, we can see this for “Belgrade Invasion”, an event that describes a part of the NATO bombing of Yugoslavia (represented in Wikidata by Q155723) but is not represented on its own. Additionally, each of those labels can be seen as a pointer to a narrative describing the event as denoted by $\eta(e, n)$. The η function returns the narrative (i.e., a directed graph as described in (1)) that further describes e , as seen for “Redivision of the World” and “Iraq War” in our running example. Additionally, this is illustrated in Fig. 2 with events e_1 and e_{11} . We call nodes, for which η yields a narrative a *recursive node*. For non-recursive nodes ne we set $\eta(ne, n) = \emptyset$.

Furthermore, we introduce *bindings* $\mu(e, n)$ for an event e in narrative n . While η allows us to recursively traverse an event in the

narrative by following the label to a sub-narrative, the binding μ maps the event label to its equivalent event in an ECKG. This is illustrated in Fig. 2 with bindings of events e_2 and e_{12} to a corresponding event in an ECKG. Additionally, we can observe a binding in Fig. 1, where μ (“NATO Eastward Expansion”) provides us with the equivalent event in Wikidata (i.e., entity Q112127201). That is, it returns the subgraph consisting of all triples with the identifier Q112127201 in the subject. Note, that a label may in practice refer to several ids, i.e., homonyms might be a problem for μ . We will discuss difficulties in this regard in Sect. 3.2.

As mentioned before, another key aspect of narratives is the narrator. In this paper we simplify the narrator and represent her as a viewpoint. Therefore, a narrative can differ depending on the viewpoint, e.g., regarding the events chosen to represent it. With respect to Sect. 2.2, this means that all information provided by a binding (here: all triples concerning the given event in an ECKG) must be viewpoint-compatible with the narrator.

3 MINING AND BINDING NARRATIVES

Based on the theoretical framework introduced in Sect. 2 we can now introduce concrete event relationships that may be part of narratives. We describe how narratives can be bound to ECKGs, i.e., how correspondences between events in narratives and events in ECKG can be defined. Finally, we introduce a method to construct narratives as synthesis from multiple text sources as one possible way of narrative mining.

3.1 Narrative Relationships

A key component in narratives are the narrative relationships $\mathcal{R}_N$. We can further characterize such relationships as triples $\langle ev_1, p, ev_2 \rangle$ with $ev_1, ev_2 \in \Gamma$ and $p \in \Sigma_{NR}$ with Σ_{NR} being a relation vocabulary. Of course, the specific manifestation of Σ_{NR} depends on the domain, i.e., the predicate labels used in narrative relationships must be defined by a domain expert. However, we will define characteristics of a typical narrative relationship that can be used as guidance in this process.

We ground our deliberations in related work concerning event-event relationships [15]. Please note that we separate the vocabulary for constituent relationships and narrative relationships. Although two events can form a narrative according to our definition, it is not sensible to interpret events in constituent relations as a narrative (otherwise each subevent relationship could be interpreted as narrative). Hence, progress between both events should be described in order to define two events as a narrative. In general, we define three categories of narrative relationships:

- Temporal Relationships:** This category of relationships is concerned with the temporal order of events in a narrative. Such relationships can be expressed by predicates like *before*, *after*, and *during*.
- Contingency and Causality:** The second category concerns causality and conditions for events. That is, this category comprises predicates like *caused by*, *lead to*, and *has effect*.
- Association:** The last category concerns associations, i.e., two events are associated on a general level. The corresponding labels are *association* or *associated with*.




Figure 3: Sketch of our narrative mining algorithm.

According to Def. 1, all events are associated with a time they took place, and hence, all events are in a temporal relationship with each other. However, from a narrative point of view, some events are mentioned together, but their temporal relationship does not matter in the narrative. Take, for instance, the association of the Belgrade Invasion, the Iraq War, and the War in Libya in Fig. 1. One could add temporal relationships between those three in the order mentioned before, but the temporal order does not matter in Putin’s speech. It is only necessary to connect those three events in the same context, i.e., the same recursive node.

Additionally, those categories form a hierarchy, i.e., if there is a causal relationship between two events, they are also in a specific temporal relationship and are associated. This is not true the other way around. As mentioned before, while all events are, by definition, in a temporal relationship, the choice of relationship in the narrative depends on the narrator. That is, associations are not superfluous as a category but can be used in cases where temporal and causal relationships are not known or not important.

3.2 Binding Events in Narratives to ECKGs

As discussed before, narratives can be grounded in ECKGs via a binding function μ . A bound narrative constructs a path in an ECKG by linking its event labels to knowledge graph events and connecting them by narrative relationships. Hence, a binding can be interpreted as substituting the event label in the narrative by the respective subgraph in an ECKG. If no corresponding event exists for an event label, the node in the narrative can be seen as virtual.

Binding the events mentioned in a narrative to named events in an ECKG provides two advantages. First, it allows for disambiguation of the event, i.e., it is clear to which real-world event the narrative refers. This eases the comparability between narratives since we compare knowledge graph entities instead of labels. Second, the binding contextualizes the event by embedding it into a subgraph provided by the ECKG. Furthermore, it includes additional metadata that can be used in downstream tasks like narrative comparison but may also include attributions from the ECKG as long as they are viewpoint-compatible with the narrator.

The whole process can be seen as a variant of entity-linking or the recently introduced event linking [50]. According to Def. 5 the

nodes of a narrative are either events or entities. If a node is an entity it must be linked to an event in a role relationship $r \in \mathcal{R}_R$. In this case: if the corresponding event can be bound to an ECKG it is likely, that r is also reflected in the ECKG. In the case of event labels, we have to differentiate whether it is a recursive node or not. We discuss the implications of both cases for the implementation of μ accordingly.

Non-recursive node. As mentioned in Sect. 2.3, events in a narrative are represented by labels like “Iraq War”. Specifically, we talk about event labels e with $\eta(e, n) = \emptyset$ here. We can identify three cases, a binding function needs to account for:

Direct Binding: The event label can be directly and unambiguously matched to an event in an ECKG. In this case, the label can be substituted by the event identifier of the KG and treated as synonymous with the subgraph of the event (i.e., the set of all triples with this specific event as a subject in an ECKG). For our running example in Fig. 1, this is the case for many events, e.g., the Iraq War and the NATO eastward expansion.

Indirect Binding: In some cases, a node in a narrative can not be directly matched on an event in an ECKG but a similar one. This is the case for both the Belgrade Invasion and the collapse of the Soviet Union in our running example.

No Binding: Suppose a node in a narrative represents not a named event, i.e., no corresponding event in an ECKG can be found. In that case, we substitute the label with a pointer to a subgraph that combines all known information of this event. In contrast to Direct Binding, this subgraph only exists in the narrative and is not materialized in the ECKG. For practical implementations, it might also be sensible to infer properties like the event type and a label for it.

The indirect binding case can become complicated since it requires a binding function to capture at least a part of the narrator’s intention, i.e., the label should provide enough information in that regard. For instance, in the Belgrade invasion, the mapped Wikidata entity represents the NATO bombing of Yugoslavia in 1999, which is close to what Putin wanted to express in his narrative. The collapse of the Soviet Union, however, is linked to the entity representing the formal dissolution of the Soviet Union, which does not explicitly

match the narrator’s intention. Deciding on a fitting entity in this specific case can hence be seen as a *loss of focus* problem, i.e., the decision is whether a more general entity still encodes the intended semantic or is too general.

Recursive Node. We evaluate $\eta(e, n)$ for recursive nodes and apply μ for each event in the generated narrative. Afterward, we again apply η for each event and continue recursively. The recursion terminates if $\eta(e, n') = \emptyset$ for all e in all narratives n' traversed in the process.

3.3 Towards Multi-Viewpoint Narrative Mining

As discussed in Sect. 2.3, our narratives resemble hierarchical event schemas. In a recent work [25] an incremental prompting approach was introduced to generate such schemas from a small document collection that contains instances of the respective event. This particular approach has two advantages: First, it requires only slight supervision for the schema generation (i.e., only a small document collection), and second, it is domain-independent.

In this section, we introduce an algorithm based on [25] for mining recursive narratives for a specific named event from a document collection. The basic idea is to integrate the narrator in this setting, i.e., we assume a collection of documents for a given event ev and a viewpoint v to differ in terms of their narratives concerning ev . That is, we do not try to abstract from specific scenarios to find event schema graphs but to find the main events ev comprises according to v . This step is performed recursively to generate narratives as depicted in Fig. 1. The algorithm is depicted in Fig. 3. A reference implementation and all prompts and data used in this paper are published on Github.² We discuss the algorithm step-by-step in the following.

Input and Output. The algorithm takes an event ev (e.g., the Iraq War or the Russo-Ukrainian War), a viewpoint v , and a document collection D as an input. The output of the algorithm is a narrative concerning ev from the perspective of v as reported in a subset $D_v \subseteq D$, i.e., all documents that can be seen as representative of v . The algorithm utilizes a large language model (LLM) through prompting

Document Selection. As a precursor, we determine the set D_v in this step. We follow [31] and proxy v by a set of documents that describe ev . Depending on v , D_v could, for instance, be a set of newspaper articles from media outlets typically associated with v or speeches by a specific speaker analogous to our running example. This step can be done, e.g., by pre-selection and annotation of documents with the respective v or by semi-automatically determining viewpoint clusters as illustrated in [35].

Event Detection. After D_v is determined, the actual first step of the algorithm is to detect, whether a document $d \in D_v$ actually refers to ev . We prompt the LLM to verify whether an article is clearly about ev and also provide a timespan t for ev in the prompt. The timespan is used for disambiguation purposes for opaque event labels (for instance it is easy to mix the Gulf War in 1991 and the Iraq War in 2003 in an article).

Timeline Extraction. The next step is analogous to the main event extraction in [25]. We generate a document event timeline $\tau(d)$ for each $d \in D_v$. An event timeline consists of all events in d that are

considered the main events concerning ev . This is done in three steps: (1) We use the prompt skeleton “Here is a news article: ‘{text of d }’ List the major events of the { ev } in chronological order as reported in the article. Keep it concise and remove everything unrelated.” to receive a list of events concerning ev in d . (2) We use few-shot prompting to generate a precise description for each event in the list (i.e., a headline). This label should include a description and a point in time or time span describing when the event in question happened. More concrete, we utilized few-shot examples “{event} – {event label (time)}” where *event* is a sentence describing the event and *event label* is a headline for this event including when it happened. (3) We again utilize few-shot prompting to 1. ensure that each event in the timeline is mentioned in d (like in (1)) and 2. filter all events outside the timespan t defined in Event Detection.

Pairwise Matching. To generate a narrative, we have to unify all event timelines, i.e., for all $d \in D_v$ we combine the respective document timelines $\tau(d)$ and find connected components by clustering. For this, we conduct a pairwise comparison between the events of two event timelines $\tau(d_1)$ and $\tau(d_2)$ and repeat this process for each two documents in D_v . We utilize SBERT [38] to compute the Euclidean distance between each event label and HDBSCAN to cluster event labels in order to find the main events over all documents in D_v . Therefore, the result of this step is clusters of event labels describing the same event in the narrative.

Synthesis. The last step is to synthesize the event clusters from the previous step into a chain of events. For each cluster, again, we generate a concise description, i.e., an event label on the basis of the event labels in the cluster via prompting. For instance, in Fig. 3 in the Synthesis step, events A and B have been identified to be in a cluster during the pairwise matching and are now synthesized into event label X. We additionally merge two clusters and prompt for a label if the labels have a cosine similarity below a given threshold. Finally, we prompt for a label to determine the narrative relationships between all events in the timeline.

Recursion. After constructing the narrative, we repeat the generation process for each node in it. In our example, we generate narratives for events A, B, C, and D, e.g., compute $\eta(A, X)$. The recursion stops if $\eta(ev, n) = \emptyset$ for all events in all narratives. To not lose the relationship between the original event and ev , we adapt the prompt in event detect to also include ev .

4 PROOF OF CONCEPT

In this section, we conduct a detailed case study concerning the narrative model as introduced in Sect. 2.3 and its applications, the mining algorithm introduced in Sect. 3.3, and finally, event linking to Wikidata.

4.1 Case Study and Document Collection

The case study concerns the Iraq War (2003–2011) as an event. We chose this event as an example because 1. it generated international tension, and hence multiple viewpoints on this event exist, and 2. it is already completed and well-documented (in contrast to our running example, the Russo-Ukrainian War).

As viewpoints, we chose Russia, the United Kingdom, and the USA, i.e., $\mathcal{V} = \{\text{RU}, \text{UK}, \text{US}\}$. Both UK and US were directly involved in the war as allies. In particular, the US were the driving force

²<https://github.com/fploetzky/WebSci2024>

behind starting the war in the first place. Russia was chosen because its government was skeptical of US claims and is naturally more reserved towards the United States. Hence, those three viewpoints will most likely illustrate different narratives on the Iraq War.

We designed the following research questions for our case study:

RQ1 How is the quality of mined narratives from open-domain web data concerning the event labels and recursion?

RQ2 What are the potentials of narrative bindings, i.e., to which degree can we bind events in the mined narratives to Wikidata?

RQ3 What are the commonalities and differences between narratives with respect to the three viewpoints?

Since this work introduces a novel view on events, no prior datasets exist that fit our needs to the best of our knowledge. Hence, we collected a small dataset of newspaper articles concerning the Iraq War from well-known news outlets for each viewpoint. This dataset is used as a proxy for the respective viewpoints. Tab. 1 depicts an overview regarding the outlets chosen and how many articles were considered. That is, for instance, for viewpoint RU we collected 197 documents from two news outlets. We chose the outlets based on their popularity scores in Media Cloud collections [39] and whether an English version of the articles is available. Otherwise, the outlet choice was random since the mining of narratives should be an abstraction of the underlying sources regardless of their structure.

Unfortunately, searching in Media Cloud for fitting articles was problematic since both the search interface and API have been under maintenance during the writing of this paper. Instead, we used the Boolean query “iraq war 2003” on DuckDuckGo to find article candidates on the respective newspaper homepages. For each viewpoint we considered the first seven DuckDuckGo result pages. Additionally, we manually discarded articles whose URLs contained obvious references to other events as false positives (this was often the case in Russian newspapers that draw inferences from the Iraq War to the ongoing Russo-Ukrainian War). Of course, this selection process could be more optimal and smaller in scale. However, it shows the difficulties of event extraction and mining in real-world Web scenarios like heterogeneity of data, documents that only partially concern the desired topic, and no shared vocabulary.

4.2 Mining and Binding Narratives

For this case study, we applied our mining algorithm to the document collections of all viewpoints. All documents were pre-processed with newspaper3k³ to extract the title and body text. To prevent articles that were too long or too short, we only considered articles between 1000 and 8000 characters. We set the timespan to include all events between 2000 and 2012 so as not to miss potential war starters and effects that happened after the war. After those two steps we ended with a total of 154 documents for RU, 188 documents for UK, and 207 documents for US. As a language model we utilized Llama-2-70b-chat-hf⁴ for this study since it is a freely available state-of-the-art LLM that runs on local hardware (in contrast to GPT3 used in [25]). Regarding SBERT we used all-mpnet-base-v2

Table 1: Document count statistics

Viewpoint	Media Outlet	# Documents
RU	sputnikglobe.com	59
	rt.com	138
UK	theguardian.com	158
	independent.co.uk	128
US	latimes.com	145
	washingtonpost.com	59
	nytimes.com	158

as model. HDBSCAN was utilized with an ϵ of 0.1, a minimum cluster size of 10, and a minimum of 2 samples after experimentation with different parameter sets. Due to the rather small number of documents, we showcase the narrative and the first recursion level. All experiments have been conducted on a server with 8 NVIDIA A40 GPUs, 2TB RAM, and an Intel(R) Xeon(R) Gold 6336Y CPU with 96 cores. As an illustration, a portion of those three narratives, along with their bindings, is depicted in Fig. 4 and will be the basis of further discussion. The colors surrounding the event labels indicate whether this event label was found in the respective viewpoints. For readability, we unified the label if they occurred in multiple viewpoints. We discuss **RQ1** with respect to the event labels, recursion, and narrative relationships separately. Afterward, we evaluate **RQ2** accordingly.

Nodes and Event Labels. After applying our algorithm to the dataset described above, we ended up with three narratives comprising 16 event labels for US, 10 event labels for UK, and 7 for RU. This reflects the size of D_v for the three viewpoints as depicted in Tab. 1. The first recursion level totaled in 28 events for US, 35 for UK, and 9 for RU; a major increase for UK due to one of the events in the UK narrative being the Iraq invasion. Therefore, in total our mining algorithm generated 105 event labels.

A major obstacle for our mining algorithm is the concise description of events. We explicitly decoupled mining and binding events to allow the algorithm to be as domain-independent as possible. This also means that no shared vocabulary exists for the event labels, and hence, this can be seen as a variant of open domain event extraction [48]. However, besides this difficulty, the quality of the event labels turned out to be good when it comes to the actual label and worthy of improvement when it comes to the extracted time intervals. Regarding the quality of the event labels, consider “Hutton Inquiry Rules Dr. Kelly’s Death a Suicide” and “Abu Ghraib Prison Abuse Scandal” in Fig. 4 as representatives of all other labels.

However, as mentioned, when it comes to the time of an event, was is oftentimes not possible to extract an exact time, or the time was simply wrong. An example of the latter is “May 2017: High Court Rules in Favor of Prosecution in ‘R v Jones’ Case, UK Military Leaders Consider War Crimes Immunity for Iraq Invasion during the Iraq War” from the UK viewpoint. While the long label is correct and the event covers the Iraq War, the court case took place in 2006 and not 2017 [28]. Concerning the connotation of the Iraq War, we only observed one event label out of all 105 events that did not

³<https://newspaper.readthedocs.io/en/latest/>

⁴<https://huggingface.co/meta-llama/Llama-2-70b-chat-hf>




Figure 4: An excerpt of a narrative mined by our algorithm.

belong to the Iraq War, namely the Russian annexation of Ukraine that was mentioned in the UK in one instance.

Narrative Relationships. In this proof of concept, we limited ourselves to one temporal relationship, specifically “happened after”. However, generating those labels proved to be the most inconsistent part of the algorithm. First, the results of prompting the LLM to generate event timestamps were only partially successful and oftentimes erroneous, as mentioned before. Future approaches may, therefore, utilize external tools for this task. Second, even with correct time annotations in the event labels, drawing the correct “happened after” relationship turned out to be a challenging task in our algorithm for the LLM.

Recursion. The results of the recursion step depended heavily on the event label generated during the first run of the algorithm. For instance, for the UK, one of the events mined for the Iraq War was the Invasion of Iraq as depicted in Fig. 4. For this event alone, 15 events have been created during the recursion step. That is, a more general label will naturally produce narratives with more events. In some cases, the general label is just refined. For instance, the narrative for the Abu Ghraib prison abuse scandal in Fig. 4 includes an event “US Military Personnel Convicted for Abu Ghraib Prison Abuse, 2004”.

In summary, for **RQ1**, we conclude that the event abstraction and recursion parts in our mining algorithm can produce hierarchical narratives. Bindings and the utility of such narratives are studied in the remaining research questions. The main obstacle here is the extraction of event times and the narrative relationships. For this study, we limited ourselves to only a single relationship, but even this rather simple temporal relationship proved to be problematic. Improvements should hence focus on this part first.

Event Linking to Wikidata. For **RQ2**, we studied the binding of events in our narratives to Wikidata as our ECKG choice. We have chosen Wikidata due to its scale and popularity as the largest freely available knowledge graph. We did not implement a binding algorithm for this proof of concept. Instead, we focused on detecting potential problems such an algorithm would face with respect to Sect. 3.2 and recent event linking literature [29, 50]. For the bindings, we used Wikipedia as a proxy to search for titles similar to the event label we tried to bind. That is, we used the event label generated by our algorithm and searched for a fitting Wikipedia article title or chapter title in an article with a title closely related to the event label. The latter was done to cater to the aforementioned indirect binding case, e.g., to find the Belgrade Invasion in the Wikipedia article concerning the NATO bombing of Yugoslavia in 1999. We chose Wikipedia as a proxy because 1. recent works in event linking identified Wikipedia as a suitable linking destination [50], and 2. Wikipedia articles often have a corresponding entity in Wikidata that shares the same title. Fig. 4 annotates the Wikidata identifier for all events for which binding was possible. In summary, we have been able to bind 22 of the 44 events in US, 22 out of 45 for UK, and 6 out of 16 for RU (i.e., 50 out of 105 in total) to Wikidata. Note that this can not be seen as a binding quality indicator since, naturally, not all events in a narrative will have a counterpart in an ECKG, as discussed before. We identified, however, several problems, in particular: (1) compound event labels, i.e., event labels that are actually referring to two events and should hence be split. (2) Double matches in bindings, for instance, “US and Iraqi Forces Launch Offensives in Fallujah, 2004” which can refer to two events in Wikidata since there have been two battles in Fallujah in 2004 (i.e., the timestamp 2004 is too general). (3) The

indirect binding case as discussed in Sect. 3.2 was observed multiple times. For instance, for “UN Issues Iraq Ultimatum: Disarm or Face Military Action (2002)” where a possible binding could be the whole relationship of the UN security council regarding the Iraq War (Q7888917). However, this does not only include the Iraq Ultimatum but also different events like the Blix investigations. The opposite case can be observed for the recursive node “Invasion of Iraq” in Fig. 4, which should refer to the War in Iraq (Q545449) and not only the actual invasion (Q107802) when considering the narrative generated from this node (the Abu Ghraib prison abuse scandal happened after the invasion but in the Iraq War). Hence, binding a recursive node requires an understanding of its inner structures. Regarding our previous assumption of connecting events that are otherwise unconnected in ECKGs, this turns out to be true for some events. For instance, Colin Powell’s address is represented in Wikidata (Q1639325) but is not connected to the Iraq War. Even if it was connected via a part of the relationship, the prominent role this event played for some viewpoints is only reflected in the narrative.

We conclude the discussion of **RQ2** and remark that binding a recursive narrative to an ECKG is a non-trivial task that requires in-depth analysis in future work.

4.3 Narrative Comparison

In this section, we discuss **RQ3** and compare the different narratives generated for this case study. Additionally, we hint at future applications for our narrative model.

Commonalities. As hypothesized in the introduction, different narratives of an event will highlight different aspects of it. Before focusing on the difference, we first investigated which events are shared by all three narratives. Examples include the actual invasion of Iraq and the Abu Ghraib prison scandal. However, the latter was featured in the main narrative in US and RU, but only after the first recursion in UK (in the invasion of Iraq event; hence, it is located there in Fig. 4). In some cases, the narratives shared the same event but with different labels, e.g., George W. Bush’s “Mission Accomplished” speech in US and UK (overlapping in Fig. 4). Also, all narratives mention at some point the WMDs that were allegedly possessed by the Iraqi government.

Differences. We start the differences with exactly those WMDs. While in RU and US they are introduced in Colin Powell’s address to the UN security council (analogous to Putin’s argumentation in Fig. 1), UK media highlights the WMD claim made in a document from the British government (wrongly dated by the algorithm). Overall, the UK viewpoint focused more often on British entities, especially former British Prime Minister Tony Blair. Note that this can be a sample bias due to the small document collection and only two outlets, but it was a noteworthy observation. Hence, the model identifies the core actors and their role in the narrative.

Another main difference is the start of the narrative, i.e., the initial reason for this war. For US and UK, the 9/11 attacks have mostly been mentioned as the start of the war, while RU mentions the Colin Powell speech as its start. Of course, this can not be generalized in this small proof of concept, but it at least indicates a framing (i.e., reaction to a terror attack versus a false assumption as a war start). Following this observation allows us to derive different paths a story took from a specific point of view.

Finally, and most notably, the narratives generated during the recursion steps showed the potential to include specific events and frames that are unique to a specific narrator. We highlight two examples. First, Fig. 4 contains a narrative concerning the Hutton Inquiry for UK. This inquiry was conducted regarding the death of Dr. David Kelly, who provided anonymous information to the BBC regarding the overstatement of the dangers of WMDs as published in the government document mentioned above. Since this case is not well-known outside of UK it makes sense that it is not included in the other two viewpoints. Second, RU features an event of the US-led coalition using depleted uranium ammunition in Iraq (visible in Fig. 4). The narrative behind this event actually connects the usage of depleted uranium ammunition with an increase in congenital birth defects in Fallujah (after the aforementioned invasion). After researching our corpus, the connection between the usage of depleted uranium and birth defects is actually available in Russian media and was correctly identified and categorized by our algorithm. Capturing this small narratives and connecting them to a greater one, i.e., categorizing them, can be seen as the major future direction for utilizing narratives for complex events.

4.4 Limitations

In this section, we briefly discuss some significant limitations of our approach that can not be evaluated in this proof of concept.

Hallucinations and Stochasticity. Hallucinations are always a problem when dealing with LLMs (cf. [34]). We did not observe any hallucinations in our scenario beyond wrong dates for events (i.e., all events have been mentioned in at least one document of the respective viewpoint), but this observation can not be generalized. Additionally, due to the stochastic nature of LLMs, narrative mining is not guaranteed to yield the same graphs. In our small-scale tests, only minor changes in event labels occurred, but on a larger scale, this must be evaluated.

Scalability. In this paper we focussed on adopting [25] for our narrative mining algorithm. The main advantage here is the generation of narratives without any setup or preprocessing except by providing the documents. However, this method does not scale in general. The generation of the US graph and the first recursion step takes about 5 hours on a high-end server and uses parallelization and quantization to perform the prompting. Hence, the algorithm must be considered preliminary work but needs additions in future work for actual productive usage.

5 RELATED WORK

Computational Narratives. Research on utilizing narratives as a method for information representation and understanding was conducted in different variations over the last years (cf. [36] for a survey). However, what a narrative actually *is* varies in-between literature [16]. Hence, various models of narratives have been discussed, focussing on different narrative aspects. Connecting-the-dots techniques for constructing storylines out of news headlines [26, 42, 49] are similar to this work in their notion of event chains. In contrast to our work, however, hierarchical aspects, i.e., different levels of granularity, are not discussed. Additionally, our work is focused on whole documents and not only headlines. Utilizing multiple viewpoints for narratives have also been modeled in previous

works [19] but to best of our knowledge never used to compare different narratives and link them to knowledge graphs. The last category of related work concerns applications that utilize narratives, for instance, for news retrieval [45] or social media analysis [8]. In contrast to this paper, both categories focus on concrete tasks and use narrow definitions of narratives tailored for those.

Event Models and Schemas. Modeling events have been a topic for at least 40 years [7]. They have been modeled as simple schemas (e.g. ACE), cubes [24], or by utilizing first order logic [1, 13]. Additionally, event models like SEM [44], the F-Model [41], LODE [43], and FARO [37] have been proposed to model various aspects of events either for specific domains or in general (cf. [30] for a recent comparison of such event ontologies). Another related topic in this regard is event schema induction, i.e., the task of learning a typical schema for an event label in an unsupervised fashion. In the past, different approaches have been investigated either to learn event role-slot templates [6] or chains of events [4, 5]. Recently, the notion of event schema graphs has been proposed [23] that add the idea of connecting relationships between events and participants of events. Induction methods for such graphs have been proposed by utilizing specialized neural architectures [22, 23] or by prompting techniques for large language models such as GPT3 [25].

Event Linking and Grounding. Since events are modeled as entities in ECKGs, techniques applied in entity linking tasks can also be used for event linking. However, some works are specifically dedicated to linking events, for instance to Wikipedia [50] or Wikidata [29]. The latter work specifically considers hierarchical event grounding, i.e., event linking to complex events in Wikidata on different levels of granularity.

6 CONCLUSION

This paper introduced narratives as rich event semantics for complex real-world events. This includes connecting events to represent the inner structures of more complex events and considering multiple points of view. We argue that it becomes increasingly important to consider multiple viewpoints on real-world events in our current polarizing times. This is especially true since we observe the world primarily through intermediaries' lenses, i.e., by different narratives from various narrators. Therefore, in this paper, we unified the concepts of viewpoints, narratives, event-centric knowledge graphs into a conceptual model.

Connecting narratives mined from unstructured discourse to structured knowledge repositories is one way of achieving multiple viewpoints on events in knowledge graphs. Based on the results of our proof of concept, we argue that the recent advances in natural language processing allow us to mine narratives from different sources in acceptable quality. However, this paper can only be considered as a starting point. Further research is necessary, especially concerning the practical usage of the proposed mining algorithm and binding events in narratives to ECKGs.

ACKNOWLEDGMENTS

Supported by the Leibniz-ScienceCampus Postdigital Participation funded by the Leibniz Association (Leibniz-Gemeinschaft).

REFERENCES

[1] João Paulo A. Almeida, Ricardo de Almeida Falbo, and Giancarlo Guizzardi. 2019. Events as Entities in Ontology-Driven Conceptual Modeling. In *Int. Conf.*

- on Conceptual Modeling (ER)*. Springer, Cham, Switzerland, 469–483. https://doi.org/10.1007/978-3-030-33223-5_39
- [2] Christian Baden and Keren Tenenboim-Weinblatt. 2018. Viewpoint, Testimony, Action: How journalists reposition source frames within news frames. *Journalism Studies* 19, 1 (2018), 143–161.
- [3] Julian Borger. 2021. *Colin Powell's UN speech: a decisive moment in undermining US credibility*. The Guardian. <https://www.theguardian.com/us-news/2021/oct/18/colin-powell-un-security-council-iraq>
- [4] Nathanael Chambers and Daniel Jurafsky. 2008. Unsupervised Learning of Narrative Event Chains. In *Proc. of the 46th Annual Meeting of the ACL (ACL)*. ACL, Kerrville, TX, USA, 789–797. <https://aclanthology.org/P08-1090/>
- [5] Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised Learning of Narrative Schemas and their Participants. In *Proc. of the 47th Annual Meeting of the ACL (ACL)*. ACL, Kerrville, TX, USA, 602–610. <https://aclanthology.org/P09-1068/>
- [6] Nathanael Chambers and Dan Jurafsky. 2011. Template-Based Information Extraction without the Templates. In *Proc. of the 49th Annual Meeting of the ACL (ACL)*. ACL, Kerrville, TX, USA, 976–986. <https://aclanthology.org/P11-1098/>
- [7] Xinhong Chen and Qing Li. 2020. Event modeling and mining: a long journey toward explainable events. *VLDB J.* 29, 1 (2020), 459–482. <https://doi.org/10.1007/s00778-019-00545-0>
- [8] Saloni Dash, Arshia Arya, Sukhmidh Kaur, and Joyojeet Pal. 2022. Narrative Building in Propaganda Networks on Indian Twitter. In *Web Science Conf. (WebSci)*. ACM, New York, NY, USA, 239–244. <https://doi.org/10.1145/3501247.3531581>
- [9] Max Fisher. 2022. *Putin's Case for War, Annotated*. The New York Times. <https://www.nytimes.com/2022/02/24/world/europe/putin-ukraine-speech.html>
- [10] Lea Frermann, Jiatong Li, Shima Khanehgar, and Gosia Mikolajczak. 2023. Conflicts, Villains, Resolutions: Towards models of Narrative Media Framing. In *Proc. of the 61st Annual Meeting of the ACL (Vol. 1: Long Papers) (ACL)*. ACL, Kerrville, TX, USA, 8712–8732. <https://doi.org/10.18653/v1/2023.acl-long.486>
- [11] Simon Gottschalk and Elena Demidova. 2018. EventKG: A Multilingual Event-Centric Temporal Knowledge Graph. In *Extended Semantic Web Conf. (ESWC)*. Springer, Cham, Switzerland, 272–287. https://doi.org/10.1007/978-3-319-93417-4_18
- [12] Simon Gottschalk, Endri Kacupaj, Sara Abdollahi, Diego Alves, Gabriel Amaral, Elisavet Koutsiana, Tin Kuculo, Daniela Major, Caio Mello, Gullal Cheema, Abdul Sittar, Golsa Tahmasebzadeh, and Gaurish Thakkar. 2021. OEKG: The Open Event Knowledge Graph. In *Int. Workshop on Cross-lingual Event-centric Open Analytics co-located with The Web Conference (CLEOPATRA@WWW) (CEUR Workshop Proceedings)*. CEUR-WS.org, Aachen, Germany, 15 pages. <https://ceur-ws.org/Vol-2829/paper5.pdf>
- [13] Giancarlo Guizzardi, Gerd Wagner, Ricardo de Almeida Falbo, Renata S. S. Guizzardi, and João Paulo A. Almeida. 2013. Towards Ontological Foundations for the Conceptual Modeling of Events. In *Int. Conf. on Conceptual Modeling (ER)*. Springer, Cham, Switzerland, 327–341. https://doi.org/10.1007/978-3-642-41924-9_27
- [14] Marc Hauser, Noam Chomsky, and Tecumseh Fitch. 2002. The faculty of language: what is it, who has it, and how did it evolve? *Science* 298, 5598 (2002), 1569–1579. <https://doi.org/10.1126/science.298.5598.1569>
- [15] Yu Hong, Tongtao Zhang, Tim O'Gorman, Sharone Horowitz-Hendler, Heng Ji, and Martha Palmer. 2016. Building a Cross-document Event-Event Relation Corpus. In *The 10th Linguistic Annotation Workshop held in conjunction with the Annual Meeting of the ACL (LAW@ACL)*. ACL, Kerrville, TX, USA, 1–6. <https://doi.org/10.18653/V1/W16-1701>
- [16] Brian Felipe Keith Norambuena, Tanushree Mitra, and Chris North. 2023. A Survey on Event-Based News Narrative Extraction. *ACM Comput. Surv.* 55, 14s, Article 300 (jul 2023), 39 pages. <https://doi.org/10.1145/3584741>
- [17] Graham Klyne, Jeremy Carroll, and Brian McBride. 2014. RDF 1.1 Concepts and Abstract Syntax. <https://www.w3.org/TR/rdf11-concepts/>
- [18] Hermann Kroll, Denis Nagel, and Wolf-Tilo Balke. 2020. Modeling Narrative Structures in Logical Overlays on Top of Knowledge Repositories. In *Int. Conf. on Conceptual Modeling (ER)*. Springer, Cham, Switzerland, 250–260. https://doi.org/10.1007/978-3-030-62522-1_18
- [19] George Lakoff and Srinu Narayanan. 2010. Toward a Computational Model of Narrative. In *Comp. Models of Narrative, AAAI Fall Symposium (AAAI Technical Report, Vol. FS-10-04)*. AAAI, Washington DC, USA, 21–28. <https://aaai.org/papers/02323-2323-toward-a-computational-model-of-narrative/>
- [20] Gregor Leban, Blaz Fortuna, Janez Brank, and Marko Grobelnik. 2014. Event registry: learning about world events from news. In *Int. World Wide Web Conf. (WWW)*. ACM, New York, NY, USA, 107–110. <https://doi.org/10.1145/2567948.2577024>
- [21] Kaley Leetaru and Philip Schrodt. 2013. *GDELT: Global data on events, location, and tone, 1979–2012*. Technical Report. Pennsylvania State University.
- [22] Manling Li, Sha Li, Zhenhailong Wang, Lifu Huang, Kyunghyun Cho, Heng Ji, Jiawei Han, and Clare R. Voss. 2021. The Future is not One-dimensional: Complex Event Schema Induction by Graph Modeling for Event Prediction. In *Conf. on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, Kerrville, TX, USA, 5203–5215. <https://doi.org/10.18653/v1/2021.emnlp-main.422>

- [23] Manling Li, Qi Zeng, Ying Lin, Kyunghyun Cho, Heng Ji, Jonathan May, Nathanael Chambers, and Clare Voss. 2020. Connecting the Dots: Event Graph Schema Induction with Path Language Modeling. In *Conf. on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, Kerrville, TX, USA, 684–695. <https://doi.org/10.18653/v1/2020.emnlp-main.50>
- [24] Qing Li, Yun Ma, and Zhenguo Yang. 2017. Event Cube – A Conceptual Framework for Event Modeling and Analysis. In *Int. Conf. on Web Inf. Systems Engineering (WISE)*. Springer, Cham, Switzerland, 499–515. https://doi.org/10.1007/978-3-319-68783-4_34
- [25] Sha Li, Ruining Zhao, Manling Li, Heng Ji, Chris Callison-Burch, and Jiawei Han. 2023. Open-Domain Hierarchical Event Schema Induction by Incremental Prompting and Verification. In *Proc. of the 61st Annual Meeting of the ACL (Vol. 1: Long Papers) (ACL)*. ACL, Kerrville, TX, USA, 5677–5697. <https://doi.org/10.18653/v1/2023.acl-long.312>
- [26] Brian Keith Norambuena and Tanushree Mitra. 2020. Narrative Maps: An Algorithmic Approach to Represent and Extract Information Narratives. *Proc. ACM Hum. Comput. Interact.* 4, CSCW3 (2020), 1–33. <https://doi.org/10.1145/3432927>
- [27] Andreas Opdahl and Bjørnar Tessem. 2021. Ontologies for finding journalistic angles. *Softw. Syst. Model.* 20, 1 (2021), 71–87. <https://doi.org/10.1007/s10270-020-00801-w>
- [28] Samuel Osborne. 2017. *High Court blocks bid to prosecute Tony Blair over Iraq War*. Independent. <https://www.independent.co.uk/news/uk/home-news/tony-blair-iraq-war-high-court-bid-prosecute-family-soldiers-killed-relatives-a7868451.html>
- [29] Jiefu Ou, Adithya Pratapa, Rishubh Gupta, and Teruko Mitamura. 2023. Hierarchical Event Grounding. In *AAAI Conf. on Artificial Intelligence (AAAI)*. AAAI Press, Washington DC, USA, 13437–13445. <https://doi.org/10.1609/AAAI.V37i11.26576>
- [30] Rajesh Piriyani, Nathalie Aussenac-Gilles, and Nathalie Hernandez. 2023. Comprehensive Survey on Ontologies about Event. In *Joint Proc. of the ESWC Workshops and Tutorials co-located with ESWC*. CEUR-WS.org, Aachen, Germany, 15 pages. https://ceur-ws.org/Vol-3443/ESWC_2023_SEMMES_XPEvent.pdf
- [31] Florian Plötzky and Wolf-Tilo Balke. 2022. It's the Same Old Story! Enriching Event-Centric Knowledge Graphs by Narrative Aspects. In *Web Science Conf. (WebSci)*. ACM, New York, NY, USA, 34–43. <https://doi.org/10.1145/3501247.3531565>
- [32] Florian Plötzky, Katarina Britz, and Wolf-Tilo Balke. 2023. Shards of Knowledge – Modeling Attributions for Event-Centric Knowledge Graphs. In *Int. Conf. on Conceptual Modeling (ER)*. Springer, Cham, Switzerland, 259–276. https://doi.org/10.1007/978-3-031-47262-6_14
- [33] Robert Porzel, Mihai Pomarlan, Laura Spillner, John Bateman, Thomas Mildner, and Carlo Santagiustina. 2022. Narrativizing Knowledge Graphs. In *Int. Workshop on AI Techn. for Leg. Docs. and Int. Workshop on Know. Graph Summ. co-located with the Int. Sem. Web Conf. (AI4LEGAL/KGSum@ISWC)*. CEUR-WS.org, Aachen, Germany, 12 pages. <https://ceur-ws.org/Vol-3257/paper11.pdf>
- [34] David Pride, Matteo Cancellieri, and Petr Knöth. 2023. CORE-GPT: Combining Open Access Research and Large Language Models for Credible, Trustworthy Question Answering. In *Int. Conf. on Theory and Practice of Digital Libraries (TPDL)*. Springer, Cham, Switzerland, 146–159. https://doi.org/10.1007/978-3-031-43849-3_13
- [35] Mainul Quraishi, Pavlos Falafios, and Eelco Herder. 2018. Viewpoint Discovery and Understanding in Social Networks. In *Web Science Conf. (WebSci)*. ACM, New York, NY, USA, 47–56. <https://doi.org/10.1145/3201064.3201076>
- [36] Priyanka Ranade, Sanorita Dey, Anupam Joshi, and Tim Finin. 2022. Computational Understanding of Narratives: A Survey. *IEEE Access* 10 (2022), 101575–101594. <https://doi.org/10.1109/ACCESS.2022.3205314>
- [37] Youssra Rebboud, Pasquale Lisena, and Raphaël Troncy. 2022. Beyond Causality: Representing Event Relations in Knowledge Graphs. In *Int. Conf. on Knowledge Engineering and Knowledge Management (EKAW)*. Springer, Cham, Switzerland, 121–135. https://doi.org/10.1007/978-3-031-17105-5_9
- [38] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Conf. on Empirical Methods in Natural Language Processing and Int. Joint Conf. on NLP (EMNLP-IJCNLP)*. ACL, Kerrville, TX, USA, 3980–3990. <https://doi.org/10.18653/V1/D19-1410>
- [39] Hal Roberts, Rahul Bhargava, Linas Valiukas, Dennis Jen, Momin M. Malik, Cindy Bishop, Emily Ndulue, Aashka Dave, Justin Clark, Bruce Etling, Robert Faris, Anushka Shah, Jasmin Rubinovitz, Alexis Hope, Catherine D'Ignazio, Fernando Bernejo, Yochai Benkler, and Ethan Zuckerman. 2021. Media Cloud: Massive Open Source Collection of Global News on the Open Web. In *The Fifteenth Int. AAAI Conf. on Web and Social Media (ICWSM)*. AAAI Press, Washington DC, USA, 1034–1045. <https://doi.org/10.1609/icwsml.v15i1.18127>
- [40] Charlotte Rudnik, Thibault Ehrhart, Olivier Ferret, Denis Teyssou, Raphaël Troncy, and Xavier Tannier. 2019. Searching News Articles Using an Event Knowledge Graph Leveraged by Wikidata. In *Companion of World Wide Web Conf. (WWW)*. ACM, New York, NY, USA, 1232–1239. <https://doi.org/10.1145/3308560.3316761>
- [41] Ansgar Scherp, Thomas Franz, Carsten Saathoff, and Steffen Staab. 2009. F—a model of events based on the foundational ontology dolce+DnS ultralight. In *Int. Conf. on Knowledge Capture (K-CAP)*. ACM, New York, NY, USA, 137–144. <https://doi.org/10.1145/1597735.1597760>
- [42] Dafna Shahaf and Carlos Guestrin. 2011. Connecting the Dots between News Articles. In *Int. Joint Conf. on Art. Int. (IJCAI)*. IJCAI/AAAI, Washington DC, USA, 2734–2739. <https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-455>
- [43] Ryan Shaw, Raphaël Troncy, and Lynda Hardman. 2009. LODE: Linking Open Descriptions of Events. In *Asian Semantic Web Conf. (ASWC)*. Springer, Berlin, Germany, 153–167. https://doi.org/10.1007/978-3-642-10871-6_11
- [44] Willem Robert van Hage, Véronique Malaisé, Roxane Segers, Laura Hollink, and Guus Schreiber. 2011. Design and use of the Simple Event Model (SEM). *J. Web Semant.* 9, 2 (2011), 128–136. <https://doi.org/10.1016/j.websem.2011.03.003>
- [45] Nikos Voskarides, Edgar Meij, Sabrina Sauer, and Maarten de Rijke. 2021. News Article Retrieval in Context for Event-centric Narrative Creation. In *Int. Conf. on the Theory of Information Retrieval (ICTIR)*. ACM, New York, NY, USA, 103–112. <https://doi.org/10.1145/3471158.3472247>
- [46] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (2014), 78–85. <https://doi.org/10.1145/2629489>
- [47] Tom Wilson, Kaitlyn Zhou, and Kate Starbird. 2018. Assembling Strategic Narratives: Information Operations as Collaborative Work within an Online Community. *Proc. ACM Hum. Comput. Interact.* 2, CSCW (2018), 183:1–183:26. <https://doi.org/10.1145/3274452>
- [48] Wei Xiang and Bang Wang. 2019. A Survey of Event Extraction from Text. *IEEE Access* 7 (2019), 173111–173137. <https://doi.org/10.1109/ACCESS.2019.2956831>
- [49] Zhihua Yan and Xijin Tang. 2023. Narrative Graph: Telling Evolving Stories Based on Event-centric Temporal Knowledge Graph. *J. of Systems Science and Systems Engineering* 32, 2 (2023), 206–221. <https://doi.org/10.1007/s11518-023-5561-0>
- [50] Xiaodong Yu, Wenpeng Yin, Nitish Gupta, and Dan Roth. 2023. Event Linking: Grounding Event Mentions to Wikipedia. In *Conf. of the European Chapter of the ACL (EACL)*. ACL, Kerrville, TX, USA, 2671–2680. <https://doi.org/10.18653/V1/2023.EACL-MAIN.196>