

Building Narrative Structures from Knowledge Graphs^{*}

Inès Blin^{**1,2}[0000-0003-0956-9466]

¹ Sony Computer Science Laboratories, Paris, France

ines.blin@sony.com

² Computer Science, Vrije Universiteit Amsterdam, The Netherlands

Abstract. Humans constantly create narratives to provide explanations for how and why something happens. Designing systems able to build such narratives would therefore contribute to building more human-centric systems, and to support uses like decision-making processes. Here, a narrative is seen as a sequence of events. My thesis investigates how a narrative can be built computationally. Four research questions are identified: representation, construction, link prediction and evaluation. A case study on the French Revolution, based upon Wikidata and Wikipedia is presented. This prototype helps identifying the first challenges such as dynamic representation and evaluation of a narrative.

Keywords: Narratives · Semantic Web · Ontologies · Reasoning

1 Introduction

Telling stories and creating narratives is suggested to be part of what makes us human [9,23]. Indeed, such narratives encompass key capabilities like making sense of experiences and providing explanations for a series of events [9]. These capabilities come natural for humans, who can identify participants and make connections between them. Furthermore, narratives are part of the understanding process of an experience [63]. Building systems able to represent and generate narratives would thus contribute to having more human-centric systems. Indeed, such systems would emulate better human processes such as creating narratives.

Studying narratives has gained growing interest in the latest years, and this also applies to the computer science domain. The Computational Models for Narratives Workshop Series (2009-2016, [39]) first attempted to better define narratives, and to better highlight their importance. The aim of the Text2Story Workshop (2018-2021, [11]) series focused on extracting narrative structures from text input. Recently, [12] explores how to leverage streams of news data to extract events and narrative structures. It provides an overview of the state-of-the-art in event extraction, temporal and causal relations, as well as storyline extraction.

Despite the existing research on narratives, no consensus was reached on how to represent and construct them. To make their construction more manageable,

^{*}



The work reported in this paper was funded by the European MUHAI project from the Horizon 2020 research and innovation programme under grant number 951846 and the Sony Computer Science Laboratories Paris.

^{**} Early Stage Ph.D. (First Year)

the literature has proposed breaking it down into smaller tasks. One simplification is to distinguish between the *fabula*, i.e. the story as it happened, and the *narration*, i.e. one expression of the fabula [38]. Take the example of a politician X giving a speech. If you are biased in favor of X , you might say that the speech was admirable, whereas if you are biased against X , you might say that the speech was political rhetoric. These two versions would be considered two different narrations of the same initial fabula: X gave a speech.

The dual process theory in cognition [18,19,20,56] states that a thought is the result of two processes. The first one searches relevant cues among a massive memory, and the second one analyses and reasons about these clues. One possible direction for implementing systems able to reproduce these two processes is to use knowledge graphs. Processes like vector space similarity, knowledge graph retrieval or graph search can be seen as the former first process. Likewise, reasoning and building a narrative based on these facts can be mapped to the second process.

My Ph.D. will use the model of a narrative as a sequence of events. Methods and approaches described in [12] also address this event-sequence type of story. The model has one perspective, hence one narration per fabula. Describing a narrative therefore boils down to identifying, describing and linking events. Categories for such representation include people, objects and locations. The aim of my Ph.D. is to leverage knowledge graphs to automatically build structured representations of narratives. System-wise, that consists in building a system that takes a graph as input and outputs a narrative graph.

Motivations for creating narratives in the form of knowledge graphs are numerous. On the representation side, knowledge graphs permit to unify data sources and enable reasoning capabilities. On the application side, they allow for semantic navigational search and reasoning on hierarchical levels. [34] furthermore argues that knowledge graphs can bring artificial intelligence to the right level of semantics and interpretability.

Lastly, narrative graph structures have multiple usages. They can first support navigation across events. In digital humanities for instance, users would first explore the collections by looking at entities and their collections, thus creating a narrative chain of events. [7]. Second, they can help in decision-making or hypotheses generation by extracting explanations, predicting events or completing a graph. Such representations have been used in domains like the biomedical domain, where a graph representation of patients' medical events is used to detect patients with venous thromboembolism [3] or maritime transportation [15].

2 Related Work

Narratives are being more and more researched, yet there is no established benchmarks on structured narratives. This section provides an overview of tasks and techniques that are relevant for narratives, even if the term “narrative” is not necessarily used. It first discusses Natural Language Processing (NLP) tasks with a focus on understanding. It then introduces graph-related tasks to help us build the narrative: graph structure and embeddings, graph search and commonsense knowledge graphs. It finally presents tasks that explicitly mention narratives.

Natural Language Processing tasks. Machine Reading Comprehension (MRC) includes tasks that aim to assess better human-like understanding. Despite recent successes on benchmarks like SQuAD 1.1 [47], there is still a gap between MRC models and human-like comprehension. [67] surveys the tasks, metrics and datasets related to this MRC research field to address these issues. It classifies the tasks according to four attributes: type of corpus, type of questions, type of answers and source of answers. Since MRC is an NLP task, most of the datasets and outputs are textual, hence not using structured representations. However, some of the datasets explicitly use the word “narrative” or “story”, like the ROCStories [41], or have inputs that could be considered narratives, like TellMeWhy [33]. Examples of tasks include next event prediction [41], story completion [40,26,35,68] and missing element prediction [41]. There is also work on combining language models and knowledge graphs for story generation [66]. More recently, [17] argued that MRC tasks assessed more *(text, question) pair* comprehension than *text* comprehension, and proposed text-based Noun Phrase enrichment, to recover all relations between Noun Phrases through prepositions.

Graph-related tasks. Other relevant research includes more graph-related tasks. Even if most of the research does not use the term “narrative” explicitly, it can help building components of the system.

One first area of research is graph representation, as either the graph structure itself or as graph embeddings. More recent research on knowledge graphs emphasises the need to have more event-centric structures to facilitate the analysis of sequence of events. A few event-centric knowledge graphs were built from existing generic knowledge graphs [21], whereas others were built from text in a domain like the news [49] or a novel [31]. Such structures also acknowledge the fundamental need to integrate the temporal dimension of events, and temporal and dynamic knowledge are currently being more investigated. Some work focuses on encoding the time component [10,50], whereas others investigate dynamic structures or embeddings [10,14,37,51,59,64]. A few also use a temporal representation to improve tasks like question-answering [29,36], knowledge graph completion [22,30,65], or link prediction [46,69]. Furthermore, datasets specifically targeted towards temporal question answering have also been released [28].

As mentioned in Section 1, the first process in the dual theory in cognition is to search relevant cues among a massive memory. To implement this process with a knowledge graph, techniques like entity linking or subgraph extraction can be used. Indeed, such techniques involve extracting relevant elements for a specific query. On one hand, current work in entity linking include entity-oriented search [16,43], relation discovery [27,60,61] and linking text pages to graph nodes [54]. On the other hand, subgraph extraction aims to select parts of a knowledge that are most coherent to reason about a specific query [24,58].

Section 1 also mentions a second process whose aim is to reason about collected relevant cues. There has been recent significant effort to build more specific, commonsense knowledge graphs. Such efforts comprise areas like generic commonsense [8,52], actions [31,42], social interactions [53] or psychology [55].

Finally, some work focuses on how to formally represent narratives and to extract them from inputs like texts or knowledge bases. For text, [4] inspects evolving stories in news articles, [13] proposes a benchmark for extracting temporal and causal relations from text, and [31] proposes a graph-based reasoning challenge to find the criminal in a Sherlock Holmes novel. [5] surveys methods to extract processes from text in business process management and emphasises process elements identified and evaluations performed. [32] lastly uses structured narrative representations and grounds them to knowledge repositories.

3 Problem Statement

The main novelty of my work will be to use knowledge graphs to build narratives in the form of structured graphs. Each one of the research questions below will contribute to building narrative structures from knowledge graphs.

1. **RQ-I: Representation of the Narrative.** Which representation to use for a narrative and for an event? Is there a representation that generalises well over different types of narratives?
2. **RQ-II: Construction of the Narrative.** How to gradually shift from a manual data exploration to a more automated one? How to best select relevant entities for a narrative, and extract the most meaningful subgraph from an input knowledge graph? How to convert this extracted subgraph into a narrative graph whose ontology is more suited for narratives?
3. **RQ-III: Link Prediction for Narrative Building.** Two types of links are meaningful for building a narrative. The first type of links complete the current representation, and the second ones connect past and next events. Therefore, how to predict meaningful links between events? When should two nodes be linked, and provided with which explanation? What types of links should be generated, and when should an entity be added to the graph? The main challenges will be to find adequate benchmarks and to handle the dynamic and temporal aspects of the narrative representation.
4. **RQ-IV: Evaluation of the Narrative.** Which metrics should be used to assess the quality of the constructed narrative?

Advancing in one of these research questions will permit to enhance the final system. This is particularly true for the construction of the narrative (RQ-II) and the hypothesis generation for a narrative (RQ-III), since both are complementary. Indeed, constructing a better narrative enables better hypotheses, and hypothesis generation enriches the final narrative.

4 Research Methodology and Approach

A first start is to use an intuitive representation for a narrative (RQ-I). [12] makes the distinction between two levels of analysis. The first one focuses on the event level (representation of each events separately) whereas the second one focuses on the narrative level (linking events). The Simple Event Model [62] can be the basis for describing the first level: it has four core classes describing the what, who, where and when parts of an event. [12] identifies temporal and

causal links as crucial to describe event-event links, which will be the basis for the second level. Consequently, this representation can be seen as a modified version of the Five Ws: who, what, when, where, and why.

RQ-I still remains a complex challenge. Some of those elements might be missing, incomplete or even contradictory depending on the sources. Thinking how to fill those Five Ws: the who part could be objects, animals or humans, the where part a location or a historical entity, the when part a date but also a historical entity, and the why part would be a link between two different events. Each W is hereafter depicted as a narrative dimension. RQ-I should also handle the dynamic and temporal aspect of a narrative. The aim of RQ-I will therefore be to think about how to represent a narrative, and to see if it is possible to find a common representation, or if there is a need for more specificity.

I will illustrate this with the example of a coup d'état. A coup is the seizure and deposition of a government and its powers. It is considered successful if the power is held for at least seven days [44]. For an event e_1 that is a coup, representing it in a graph would be identifying the participant p_1 , the timestamp t_1 , the location l_1 and the cause e_0 , encoded as follows: $hasParticipant(e_1, p_1)$, $hasTimestamp(e_1, t_1)$, $hasLocation(e_1, l_1)$, $hasCause(e_1, e_0)$. Furthermore, if e_1 is successful, there is change of government, hence the representation is different before and after t_1 . If e_1 happened in country c_1 , encoded information could for instance be $hasGovernment(c_1, gov_{old}, t_1^-)$ and $hasGovernment(c_1, gov_{new}, t_1^+)$.

Two main components were identified for the construction of the narrative (RQ-II): collecting elements, and building the narrative graph. The first component consists in collecting relevant elements for a given narrative, e.g. all events. One technique would be finding good techniques to optimally search this graph. The second component consists in adapting the structure and completing the graph. For instance in the aforementioned coup e_1 with participant p_1 , the input knowledge graph can include $commander(e_1, p_1)$, where the wanted output is $hasParticipant(e_1, p_1)$. In that case, the narrative graph construction model should learn to map $commander$ to $hasParticipant$.

As for hypothesis generation on narratives (RQ-III), [37] makes the distinction between offline inference (called interpolation, or graph completion) and online inference (called extrapolation, or next event prediction). Two graph-related techniques were identified: graph completion and event prediction. In a narrative setting, the former would be completing and therefore enriching an existing narrative (i.e. adding nodes and edges), while the latter would be next event prediction. One hypothesis generation could be to generalise a successful coup: $outcome(e_1, success) \leftarrow ends(e_1, gov_{old}), starts(e_1, gov_{new})$. A successful coup causes the ending of a government and the beginning of a new one.

Different hierarchies of narratives can also be defined, related to RQ-I, RQ-II and RQ-III. At least a distinction between generic and instantiated narratives can be made. An instantiated narrative has only grounded variables, a more generic narrative can also have variables. This is therefore related to graph patterns that aim to identify common sub graphs in a graph [6]. Detecting patterns in narrative graphs can thus contribute to building more generic narratives, and can furthermore be used for next event prediction or graph completion.

5 Evaluation Plan

The evaluation part of the narrative was included as a whole separate question (RQ-IV), since it is a non trivial question. Indeed, it might be complicated to evaluate a narrative benchmark as a complete end-to-end task. Nevertheless, the objective of RQ-IV will be to tackle certain components of building a narrative, and evaluate them separately. RQ-IV will also attempt to define and formalise metrics for narrative understanding.

Techniques to evaluate the narrative representation (RQ-I) include ontology evaluation methods. [45] surveys methods evaluating an ontology according to its quality and correctness and using criteria like accuracy, completeness, conciseness, adaptability, clarity, computational efficiency or consistency.

As described in Section 4, the construction of the narrative (RQ-II) can be decomposed into two steps: retrieving relevant content, then building the narrative graph. If there is the ground truth of events for a given topic, metrics like precision, recall and f1 can be used. Once the event-centric graph is built, schema-correctness can also be a way to evaluate the coherence of the graph.

For link prediction to complete a graph representation (RQ-III), one way to evaluate is to complete related challenges or benchmarks. In that case, the evaluation will use the metrics defined by those tasks. For prediction tasks, metrics like Precision, Recall or F1-score are often used. For question-answering tasks, metrics could be Mean Reciprocal Rank or hits@k. Another way of evaluating the narrative will be to define beforehand measures that can assess the understanding of the narrative [57]. In that case, the aim will be to maximise or minimise those dimensions of understanding to assess the narrative output. Such dimensions could include compatibility and relevance.

6 Preliminary Results

This section describes an initial prototype for one historical narrative, the French Revolution. Its aim was to explore Wikidata and Wikipedia to build a narrative on the French Revolution. As a historic event, the events that are part of the French Revolution are breaking points: there is a before and an after. Therefore, studying this example ensures to have a series of events referenced in knowledge bases. The case study mainly focuses on RQ-I, RQ-II, and RQ-III.

6.1 A modified Simple Event Model to represent the narrative

For the representation of the narrative (RQ-I), a modified version of the Simple Event Model [62] was used. There are four core classes in this model: `sem:Event` (what), `sem:Actor` (who), `sem:Place` (where), and `sem:Time` (when). The constraints classes `sem:Role`, `sem:Temporary` and `sem:View` can respectively add information on the role of an actor, a temporal constraint or on a specific viewpoint. This model does not however permit to link different events, nor gives relations between classes of the same type. Two types of links were thus added. The first types of links are temporal or causal links between events.

Allen’s relations [2] were used for temporal links. The `wikidata:has effect` predicate was used for causal links. The second types of links are links between core classes of the same type. The predicate `dbo:alongside` was used to denote relations across participants. An example is given in Appendix B.

There are several advantages to use the modified Simple Event Model [62]. First, it is possible to include different perspectives in this model. Indeed, the `sem:View` class allows to add properties that only hold according to a certain authority, hence allowing to compare different viewpoints. Second, regarding uncertainty in the graph, it is possible in the model to add uncertainty on time intervals. Third, the core predicates of this model allow to easily separate different types of subgraphs, like temporal or causal subgraphs for instance. It therefore permits to analyse the narrative under different angles. Lastly, the types classes `sem:EventType`, `sem:RoleType`, `sem:ActorType` and `sem:PlaceType` can help identify more generic narratives to identify narrative schemes rather than instantiated narratives.

Some other aspects are however less straightforward with that model. Indeed, the Simple Event Model described above resembles more the format of a timeline with some causal links rather than a state-based representation. Therefore, the model is less flexible to represent changes over time (either changes of nodes’ attributes and new nodes or edges that are added). With this model, temporal constraints would be the way to represent those changes, since they enable properties that hold only during a certain time interval.

6.2 Gathering data from Wikidata and Wikipedia

Regarding RQ-I, the French Revolution narrative was defined as a set of events. An event here is a node in Wikidata with a path to the French Revolution node³. The paths were manually chosen based on a Wikidata exploration and are described in Table 1. These paths enabled to collect 59 events, among which 53 unique collected events and 48 unique ones with a human-readable label.

Table 1: Graph paths used to retrieve events during the French Revolution. The path (event, part of, French Revolution) reads as follows: there is a directed edge with the label part of in the graph from the event node to the French Revolution node. *wd* is the prefix for the namespace <http://www.wikidata.org/wiki/>.

#	Human-readable path	URI Path	Number of events collected
1	(French Revolution, has significant event, event)	(wd:Q6534, wd:Property:P793, ?e)	7
2	(event, part of, French Revolution)	(?e, wd:Property:P361, wd:Q6534)	48
3	(event, is instance of, ,historical country) & (event, has country, c)	(?e, wd:Property:P31, wd:Q3024240) (?e, wd:Property:P17, wd:Q142)	4

The next step was to see how much information it was possible to gather for each narrative dimension. Specifically, the main points of interest were participants, locations and dates, as well as temporal and causal links between events – equivalent to the event-level and narrative-level links depicted by [12] and mentioned in Section 4. Two main sources of data were used:

³ <https://www.wikidata.org/wiki/Q6534>

- The attributes of each node in Wikidata, i.e. outgoing predicates and edges.
- Wikipedia Infoboxes. An Infobox in Wikipedia is a table with textual properties and attributes that contains the most important information about the current page. Most interestingly, the infoboxes contain URL links to other Wikipedia pages, which can be linked again to Wikidata.

Relevant predicates in Wikidata and attributes in Infoboxes were manually selected for narrative building. A predicate was considered relevant if it was adding information on either a participant, the type of event, a timestamp, a location or a cause. Details of predicates are available in Appendix A. Table 2 shows the number and percentage of events that contain at least one information for each narrative dimension. Overall, we see that more information is retrieved regarding places, times and temporal links between events, with Wikidata having a bit more information than Wikipedia. On the other hand, Wikipedia contains more information on participants and causal links between events. Out of the 48 events retrieved in Wikidata, only 26 of them had a corresponding Infobox in Wikipedia, resulting in loss of information and lower numbers. Furthermore, the temporal links between events are artificially boosted by the “part of” predicate in Wikidata and the “partof” attribute in Wikipedia: 42 events are directly linked to the “French Revolution” with predicate “part of”, whereas 12 events are linked to the French Revolution Wikipedia page with attribute “partof”. These do not add much information, since it was one of the path described in Figure 1. Such pairs were removed for comparison, and it was found that 8 events had a temporal link for both Wikidata and Wikipedia, 8 events a temporal link for Wikidata only and 1 event a temporal link for Wikipedia only.

Table 2: Number and percentage of events that contain at least one information for each type. WD stands for Wikidata and WP for Wikipedia. $WD \cap WP$ indicates events that were retrieved both by WD and WP for a given type, $WD \setminus WP$ events that were retrieved by WD only and $WP \setminus WD$ events that were retrieved by WP only.

Type	$WD \cap WP$		$WD \setminus WP$		$WP \setminus WD$		Total		Not retrieved	
	count	perc.	count	perc.	count	perc.	count	perc.	count	perc.
who	5	10	0	0	15	31	20	42	28	58
when	25	52	17	35	0	0	42	87	6	13
where	17	35	19	40	4	8	40	83	8	17
causal link	1	2	0	0	12	25	13	27	35	73
temp. link	21	44	27	56	0	0	48	100	0	0

6.3 Building the narrative network

The final step was to construct a narrative graph from the content gathered and described in Section 6.2 (RQ-II). Figure 1 presents the steps followed to build the final output graph. As described in Section 6.2, events are first searched manually through Wikidata, and relevant paths are chosen. Using those paths, all events for the experiment are then collected, and enriched with infobox information from Wikipedia. The data collected is then converted to triples that enrich the final graph. This section describes more in details how the graph was built.

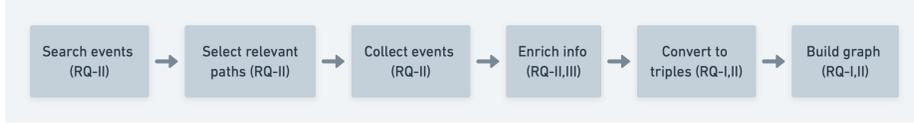


Fig. 1: Pipeline of the steps followed to build the final graph. Related research questions are also added. Data was used from Wikidata and Wikipedia.

For the construction of the narrative (RQ-II), the objective was to convert the original triples in Wikidata and the key-value pairs in Wikipedia to a format for narratives. The rules for conversion were manually designed, with the following strategy: for each event i) for each url link in the infobox, find the corresponding Wikidata URI and add the triples to the output graph ii) convert the triples (s, p, o) , with s the URI of the event and p a relevant predicate, from Wikidata.

The graph for the event 13 Vendémiaire is displayed in Figure 2. One can understand that 13 Vendémiaire was a coup d'état between Royalists and Republicans, and that Paul Barras and Napoleon were Republicans. Some limitations also appear in that representation. First, the output of this coup is missing: who was at the origin of this coup, was it successful? Second, the node “First French Republic” is overloaded with too many meanings, as it is both a place and an actor. Semantically, it should probably be considered two different entities.

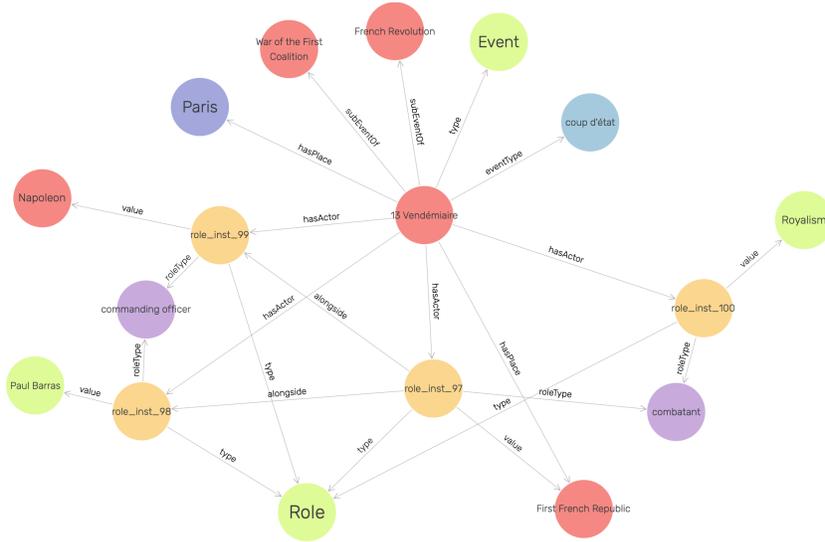


Fig. 2: Visual graph of the 13 Vendémiaire event. The `role_inst_{k}` nodes are encodings for blank nodes in the graph. One can see that the event is a coup d'état with two combatants: the First French Republic and Royalism. One therefore concludes that the event was a conflict between royalists and republicans.

For hypothesis generation on the graph (RQ-III), the main reasoning step was to manually define the rules defined above, however it does not add new knowledge with regards to the data gathered. Simple reasoning steps could be

to enrich the types of nodes added to the graph: for instance, if an event X is a sub-event of another node Y unseen so far in the graph, then Y is also of type event. To prepare for future reasoning, using labels that give us more information about semantics could also help this process.

One improvement to the manual path selection would be to have the machine learning how to collect such events automatically with a knowledge graph. A graph search experiment is currently being worked on, where the aim is to find good heuristics to explore a knowledge graph to retrieve events for a narrative.

7 Conclusions and Lessons Learnt

In this paper, I presented my Ph.D. work that aims to construct narrative networks based on knowledge graphs. It explained the research questions addressed and provided a survey of the state-of-the-art of relevant research. The prototype described in Section 6 focused on using information from Wikidata and Wikipedia to build a narrative network. The aim was to see how much information was available to describe elements like participants, locations and cautions, as well as to provide a first structured representation.

The prototype furthermore enabled to identify the first challenges for building narrative networks from knowledge graphs. First, there is the fundamental question of the temporality and the dynamicity of the graph representation. Furthermore, how to best represent temporal changes in the narrative network? Second, there is the challenge of evaluating the narrative, and assess whether a set of events is a good one to describe a narrative. How many links should be added to the graph to consider the narrative complete, or correct? Third, there is the scaling question, related to how many narratives one can build re using the same process. For instance, using the same process as for the French Revolution, how easy or hard would that be for another revolution? Lastly, there is the question of relevant input resources to use to build the narrative. Using generic knowledge graphs [25] can be a starting point, but sometimes domain-specific knowledge graphs [1,48] are more suited. In any case, it is important to remember that the final narrative structure will have be biased towards the content of the input resources.

The next steps for my Ph.D. will be to refine over the initial prototype to improve the narrative building process. Future work will especially be about automating more the components for narrative building, like searching relevant entities or completing the graph. Furthermore, the work presented on RQ-I and RQ-II focused more on the representations of events than representations of temporal relations. Future work will therefore also be on such temporal relations.

Acknowledgements I thank Annette ten Teije (VUA), Ilaria Tiddi (VUA), and Remi van Trijp (CSL) for their comments and feedbacks on this paper. I also thank Frank van Harmelen (VUA) for valuable advice and discussion. I thank David Colliaux (CSL), Michael Anslow (CSL), Martina Galletti (CSL) and Adam Dahlgren (Umeå University) for interesting discussion and feedbacks.

References

1. Abu-Salih, B.: Domain-specific knowledge graphs: A survey. *Journal of Network and Computer Applications* **185**, 103076 (2021)
2. Allen, J.F.: Maintaining knowledge about temporal intervals. *Communications of the ACM* **26**(11), 832–843 (1983)
3. Bakalara, J., Guyet, T., Dameron, O., Happe, A., Oger, E.: An extension of chronicles temporal model with taxonomies-application to epidemiological studies. In: *HEALTHINF 2021-14th International Conference on Health Informatics*. pp. 1–10 (2021)
4. Barranco, R.C., Boedihardjo, A.P., Hossain, M.S.: Analyzing evolving stories in news articles. *International Journal of Data Science and Analytics* **8**(3), 241–256 (2019)
5. Bellan, P., Dragoni, M., Ghidini, C.: Process extraction from text: state of the art and challenges for the future. *arXiv preprint arXiv:2110.03754* (2021)
6. Bloem, P., de Rooij, S.: Large-scale network motif analysis using compression. *Data Mining and Knowledge Discovery* **34**(5), 1421–1453 (2020)
7. Boer, V.d., Melgar, L., Inel, O., Ortiz, C.M., Aroyo, L., Oomen, J.: Enriching media collections for event-based exploration. In: *Research Conference on Metadata and Semantics Research*. pp. 189–201. Springer (2017)
8. Bosselut, A., Rashkin, H., Sap, M., Malaviya, C., Celikyilmaz, A., Choi, Y.: Comet: Commonsense transformers for automatic knowledge graph construction. *arXiv preprint arXiv:1906.05317* (2019)
9. Boyd, B.: *On the origin of stories: Evolution, cognition, and fiction*. Harvard University Press (2010)
10. Cai, L., Janowicz, K., Yan, B., Zhu, R., Mai, G.: Time in a box: Advancing knowledge graph completion with temporal scopes. In: *Proceedings of the 11th on Knowledge Capture Conference*. pp. 121–128 (2021)
11. Campos, R., Jorge, A., Jatowt, A., Bhatia, S.: The 3rd international workshop on narrative extraction from texts: Text2story 2020. In: *European Conference on Information Retrieval*. pp. 648–653. Springer (2020)
12. Caselli, T., Hovy, E., Palmer, M., Vossen, P.: *Computational Analysis of Storylines: Making Sense of Events*. Cambridge University Press (2021)
13. Caselli, T., Vossen, P.: The event storyline corpus: A new benchmark for causal and temporal relation extraction. In: *Proceedings of the Events and Stories in the News Workshop*. pp. 77–86 (2017)
14. Cochez, M., Ristoski, P., Ponzetto, S.P., Paulheim, H.: Biased graph walks for rdf graph embeddings. In: *Proceedings of the 7th International Conference on Web Intelligence, Mining and Semantics*. pp. 1–12 (2017)
15. Del Mondo, G., Peng, P., Gensel, J., Claramunt, C., Lu, F.: Leveraging spatio-temporal graphs and knowledge graphs: Perspectives in the field of maritime transportation. *ISPRS International Journal of Geo-Information* **10**(8), 541 (2021)
16. Devezas, J., Nunes, S.: A review of graph-based models for entity-oriented search. *SN Computer Science* **2**(6), 1–36 (2021)
17. Elazar, Y., Basmov, V., Goldberg, Y., Tsarfaty, R.: Text-based np enrichment. *arXiv preprint arXiv:2109.12085* (2021)
18. Evans, J.S.B.: Heuristic and analytic processes in reasoning. *British Journal of Psychology* **75**(4), 451–468 (1984)
19. Evans, J.S.B.: In two minds: dual-process accounts of reasoning. *Trends in cognitive sciences* **7**(10), 454–459 (2003)

20. Evans, J.S.B.: Dual-processing accounts of reasoning, judgment, and social cognition. *Annu. Rev. Psychol.* **59**, 255–278 (2008)
21. Gottschalk, S., Demidova, E.: Eventkg—the hub of event knowledge on the web—and biographical timeline generation. *Semantic Web* **10**(6), 1039–1070 (2019)
22. Gottschalk, S., Demidova, E.: Happening: Happen, predict, infer—event series completion in a knowledge graph. In: *International Semantic Web Conference*. pp. 200–218. Springer (2019)
23. Gottschall, J.: *The storytelling animal: How stories make us human*. Houghton Mifflin Harcourt (2012)
24. Han, Z., Chen, P., Ma, Y., Tresp, V.: Explainable subgraph reasoning for forecasting on temporal knowledge graphs. In: *International Conference on Learning Representations* (2020)
25. Heist, N., Hertling, S., Ringler, D., Paulheim, H.: Knowledge graphs on the web—an overview. (2020)
26. Huang, C.Y., Huang, T.H.: Semantic frame forecast. *arXiv preprint arXiv:2104.05604* (2021)
27. Hyvönen, E., Rantala, H., et al.: Knowledge-based relation discovery in cultural heritage knowledge graphs. In: *Digital Humanities in Nordic Countries Proceedings of the Digital Humanities in the Nordic Countries 4th Conference*. CEUR-WS. org (2019)
28. Jia, Z., Abujabal, A., Saha Roy, R., Strötgen, J., Weikum, G.: Tempquestions: A benchmark for temporal question answering. In: *Companion Proceedings of the The Web Conference 2018*. pp. 1057–1062 (2018)
29. Jia, Z., Pramanik, S., Saha Roy, R., Weikum, G.: Complex temporal question answering on knowledge graphs. In: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. pp. 792–802 (2021)
30. Jung, J., Jung, J., Kang, U.: Learning to walk across time for interpretable temporal knowledge graph completion. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. pp. 786–795 (2021)
31. Kawamura, T., Egami, S., Tamura, K., Hokazono, Y., Ugai, T., Koyanagi, Y., Nishino, F., Okajima, S., Murakami, K., Takamatsu, K., et al.: Report on the first knowledge graph reasoning challenge 2018. In: *Joint International Semantic Technology Conference*. pp. 18–34. Springer (2019)
32. Kroll, H., Nagel, D., Balke, W.T.: Modeling narrative structures in logical overlays on top of knowledge repositories. In: *International Conference on Conceptual Modeling*. pp. 250–260. Springer (2020)
33. Lal, Y.K., Chambers, N., Mooney, R., Balasubramanian, N.: Tellmewhy: A dataset for answering why-questions in narratives. *arXiv preprint arXiv:2106.06132* (2021)
34. Lecue, F.: On the role of knowledge graphs in explainable ai. *Semantic Web* **11**(1), 41–51 (2020)
35. Li, Z., Ding, X., Liu, T.: Constructing narrative event evolutionary graph for script event prediction. *arXiv preprint arXiv:1805.05081* (2018)
36. Li, Z., Jin, X., Guan, S., Li, W., Guo, J., Wang, Y., Cheng, X.: Search from history and reason for future: Two-stage reasoning on temporal knowledge graphs. *arXiv preprint arXiv:2106.00327* (2021)
37. Liao, S., Liang, S., Meng, Z., Zhang, Q.: Learning dynamic embeddings for temporal knowledge graphs. In: *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. pp. 535–543 (2021)
38. Meghini, C., Bartalesi, V., Metilli, D.: Representing narratives in digital libraries: The narrative ontology. *Semantic Web (Preprint)*, 1–24 (2021)

39. Miller, B., Lieto, A., Ronfard, R., Ware, S., Finlayson, M.: Proceedings of the 7th workshop on computational models of narrative. In: 7th Workshop on Computational Models of Narrative (CMN 2016). vol. 53 (2016)
40. Mori, Y., Yamane, H., Mukuta, Y., Harada, T.: Finding and generating a missing part for story completion. In: Proceedings of the The 4th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature. pp. 156–166 (2020)
41. Mostafazadeh, N., Chambers, N., He, X., Parikh, D., Batra, D., Vanderwende, L., Kohli, P., Allen, J.: A corpus and evaluation framework for deeper understanding of commonsense stories. arXiv preprint arXiv:1604.01696 (2016)
42. Narayanan, S.: Reasoning about actions in narrative understanding. In: IJCAI. vol. 99, pp. 350–357. Citeseer (1999)
43. Oza, P., Dietz, L.: Which entities are relevant for the story? In: Text2Story@ ECIR. pp. 41–48 (2021)
44. Powell, J.M., Thyne, C.L.: Global instances of coups from 1950 to 2010: A new dataset. *Journal of Peace Research* **48**(2), 249–259 (2011)
45. Raad, J., Cruz, C.: A survey on ontology evaluation methods. In: Proceedings of the International Conference on Knowledge Engineering and Ontology Development, part of the 7th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (2015)
46. Radstok, W., Chekol, M., Velegrakis, Y.: Leveraging static models for link prediction in temporal knowledge graphs. In: 2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI). pp. 1034–1041. IEEE (2021)
47. Rajpurkar, P., Zhang, J., Lopyrev, K., Liang, P.: Squad: 100,000+ questions for machine comprehension of text. arXiv preprint arXiv:1606.05250 (2016)
48. Reese, J.T., Unni, D., Callahan, T.J., Cappelletti, L., Ravanmehr, V., Carbon, S., Shefchek, K.A., Good, B.M., Balhoff, J.P., Fontana, T., et al.: Kg-covid-19: A framework to produce customized knowledge graphs for covid-19 response. *Patterns* **2**(1), 100155 (2021)
49. Rospocher, M., van Erp, M., Vossen, P., Fokkens, A., Aldabe, I., Rigau, G., Soroa, A., Ploeger, T., Bogaard, T.: Building event-centric knowledge graphs from news. *Journal of Web Semantics* **37**, 132–151 (2016)
50. Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., Bronstein, M.: Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637 (2020)
51. Rudolph, M., Blei, D.: Dynamic embeddings for language evolution. In: Proceedings of the 2018 World Wide Web Conference. pp. 1003–1011 (2018)
52. Sap, M., Le Bras, R., Allaway, E., Bhagavatula, C., Lourie, N., Rashkin, H., Roof, B., Smith, N.A., Choi, Y.: Atomic: An atlas of machine commonsense for if-then reasoning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 3027–3035 (2019)
53. Sap, M., Rashkin, H., Chen, D., LeBras, R., Choi, Y.: Socialiqa: Commonsense reasoning about social interactions. arXiv preprint arXiv:1904.09728 (2019)
54. Schäfer, B.: Exploiting DBpedia for graph-based entity linking to Wikipedia. Ph.D. thesis (2014)
55. Shu, T., Bhandwaldar, A., Gan, C., Smith, K.A., Liu, S., Gutfreund, D., Spelke, E., Tenenbaum, J.B., Ullman, T.D.: Agent: A benchmark for core psychological reasoning. arXiv preprint arXiv:2102.12321 (2021)
56. Sloman, S.A.: The empirical case for two systems of reasoning. *Psychological bulletin* **119**(1), 3 (1996)

57. Steels, L.: Conceptual foundations of human-centric ai (2022)
58. Teru, K., Denis, E., Hamilton, W.: Inductive relation prediction by subgraph reasoning. In: International Conference on Machine Learning. pp. 9448–9457. PMLR (2020)
59. Tiddi, I., Daga, E., Bastianelli, E., d’Aquin, M.: Update of time-invalid information in knowledge bases through mobile agents (2016)
60. Tiddi, I., d’Aquin, M., Motta, E.: Walking linked data: a graph traversal approach to explain clusters (2014)
61. Traverso-Ribón, I., Palma, G., Flores, A., Vidal, M.E.: Considering semantics on the discovery of relations in knowledge graphs. In: European Knowledge Acquisition Workshop. pp. 666–680. Springer (2016)
62. Van Hage, W.R., Malaisé, V., Segers, R., Hollink, L., Schreiber, G.: Design and use of the simple event model (sem). *Journal of Web Semantics* **9**(2), 128–136 (2011)
63. Vilarroya, Ó.: Somos lo que nos contamos. Cómo los relatos construyen el mundo en que vivimos., Editorial Ariel, Barcelona (2019)
64. Wewer, C., Lemmerich, F., Cochez, M.: Updating embeddings for dynamic knowledge graphs. arXiv preprint arXiv:2109.10896 (2021)
65. Xu, C., Chen, Y.Y., Nayyeri, M., Lehmann, J.: Temporal knowledge graph completion using a linear temporal regularizer and multivector embeddings. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 2569–2578 (2021)
66. Yang, X., Tiddi, I.: Creative storytelling with language models and knowledge graphs. In: CIKM (Workshops) (2020)
67. Zeng, C., Li, S., Li, Q., Hu, J., Hu, J.: A survey on machine reading comprehension—tasks, evaluation metrics and benchmark datasets. *Applied Sciences* **10**(21), 7640 (2020)
68. Zhang, M., Ye, K., Hwa, R., Kovashka, A.: Story completion with explicit modeling of commonsense knowledge. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 376–377 (2020)
69. Zhang, M., Chen, Y.: Link prediction based on graph neural networks. *Advances in Neural Information Processing Systems* **31**, 5165–5175 (2018)

A Predicates selected for each narrative dimension

Type of predicate	Predicates label
who	participant, organizer, founded by
where	country, location, coordinate location, located in the administrative territorial entity, continent
when	point in time, start time, end time, inception, dissolved, abolished or demolished date, publication date
temporal link between events	part of, followed by, replaces, replaced by, follows, time period
causal link between events	has effect

Fig. 3: Selected predicates for each of the narrative dimension in Wikidata.

Type of predicate	Predicates label
who	Participants, appointer, combatant{k}, commander{k}, commanders, deputy{k}, founder, house, leader{k}, legislature, organisers, p{k}, participants, precursor
where	Location, area, coordinates, location, place
when	Date, abolished, date, date_end, date_event, date_pre, date_start, defunct, disbanded, established, formation, founded_date, life_span, year_
temporal link between events	era, event{k}, event_end, event_pre, event_start, partof , preceded_by, succeeded_by, succession
causal link between events	Result, cause, outcome, result, territory

Fig. 4: Selected predicates for each of the narrative dimension in Wikipedia.

B Example of one event construction

Statements

instance of	Coup d'état	edit
	- 0 references	+ add reference
		+ add value
part of	French Revolution	edit
	+ 1 reference	+ add reference
		+ add value
image	 13 Vendémiaire.jpg 1,418 × 875, 324 KB media legend Подваление мятежа 13 вандемьера. Шарль Моне (Russian)	edit
	+ 1 reference	+ add reference
		+ add value
named after	13 Vendémiaire	edit
	- 0 references	+ add reference
		+ add value
country	French First Republic	edit
	- 0 references	+ add reference
		+ add value
location	Paris	edit
	+ 1 reference	+ add reference
		+ add value
coordinate location	 48°51'55"N, 2°19'57"E	edit
	- 0 references	+ add reference
		+ add value
point in time	6 October 1795 ^{region}	edit
	- 0 references	+ add reference
		+ add value
Commons category	13 Vendémiaire	edit
	+ 1 reference	+ add reference
		+ add value

(a) Wikidata page of 13 Vendémiaire

13 Vendémiaire

Part of the [French Revolution](#) and the [War of the First Coalition](#)



Napoleon Bonaparte's quelling of the Royalist revolt, 13 Vendémiaire, in front of the [Église Saint-Roch](#), Saint-Honoré Street, Paris.

Date	5 October 1795
Location	Paris, France 48.856667°N 2.350987°E﻿ / ﻿48.856667°N 2.350987°E﻿ / 48.856667; 2.350987
Result	Republican victory

Belligerents

 French Republic	 French Royalists
---	--

Commanders and leaders

 Paul Barras	 Richer de Sévigny
 Napoleon Bonaparte	

Strength

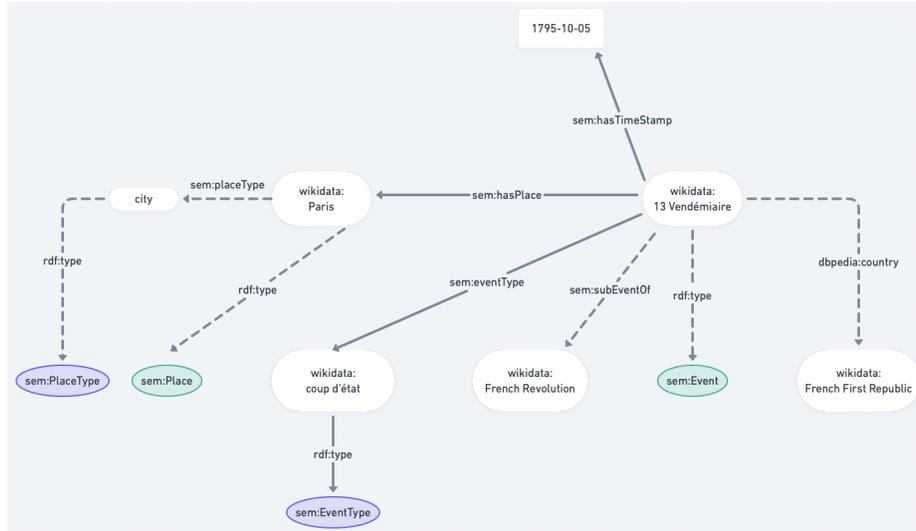
6,000 total	7,000 ^[1]
4,500 Regular Troops,	
1,500 'Patriots', 40	
cannons	

Casualties and losses

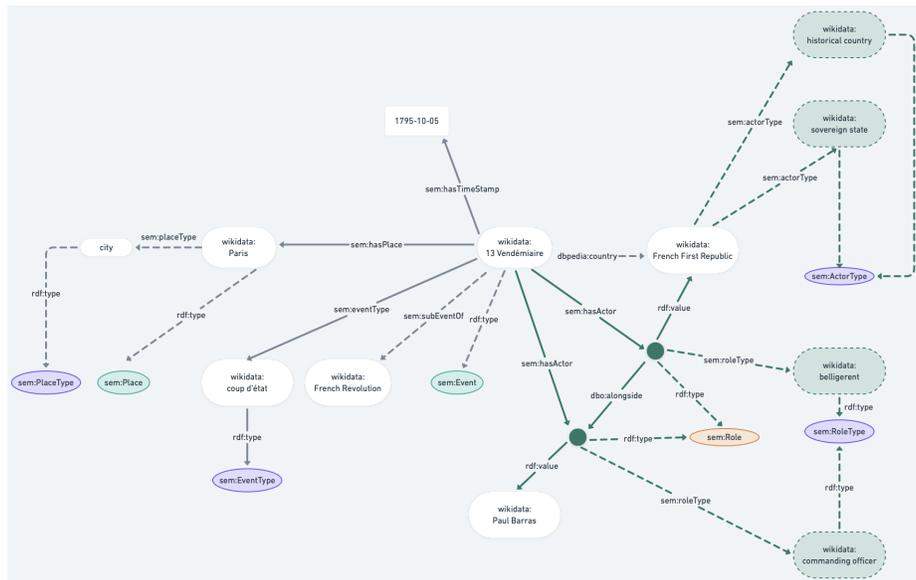
100 dead and wounded ^[2]	300 dead
	2 executed
400 civilians	

(b) Wikipedia infobox of 13 Vendémiaire

Fig. 5: Wikidata and Wikipedia page content used to build an event representation for 13 Vendémiaire. Dashed lines indicates predicates or keys that were used, and full lines values. For clarity in visualisation, not all predicates related to the narrative dimensions were used, but only a subset of them.



(a) Event representation using Wikidata useful predicates only.



(b) Event representation after both content from Wikidata and Wikipedia has been added.

Fig. 6: Event representation at different steps: using Wikidata outgoing links of the event (Figure 6a) and Wikipedia infoboxes (Figure 6b). On Figure 6b, green edges on the right indicate edges and nodes that were newly added with the Infobox. Refer to Figure from [62] for the original example.