An Ontology for Creating Hypermedia Stories Over Knowledge Graphs

Pasquale **Lisena**¹, Daniel **Schwabe**², Marieke van **Erp**³, Inger **Leemans**³,
William **Tullett**⁴, Raphaël **Troncy**¹, Sofia Collette **Ehrich**³ and Thibault **Ehrhart**¹

¹*EURECOM, Campus SophiaTech, 450 Route des Chappes, 06410 Biot, France*
²*Institute Jožef Stefan, Jamova cesta 39, 1000 Ljubljana, Slovenia*
³*KNAW Humanities cluster, Oudezijds Achterburgwal 185, 1012 DK Amsterdam, The Netherlands*
⁴*University of York, York, YO10 5DD, United Kingdom*

Abstract

Storylines are an effective and coherent mechanism to organise knowledge of a set of objects. In this paper, we present a suite of ontologies for generating storylines around cultural heritage artefacts and encyclopaedic entries. We present a use case on olfactory heritage where our ontologies are able to represent different coherent narrative structures for a set of cultural heritage artefacts.

Keywords
storylines, digital humanities, modelling cultural heritage data, navigating cultural heritage data

1. Introduction

Knowledge Graphs (KGs) have become one of the main tools used for Knowledge Management (KM) [1]. KGs are computational structures that are meant to be accessed programmatically using queries. Even so, it is also necessary for humans to consume KG contents directly. A typical approach for this has been to provide interfaces that allow users to explore the contents of KGs, for example through a faceted search. Examples of such tools, notably in the cultural heritage domain, include ADASilk [2],[1] the Odeuropa Smell Explorer,[2] CultureSampo [3],[3] and the Europeana Portal [4].[4] Storytelling has long been recognized as a knowledge management (KM) approach [5] and is used in many different settings to disseminate information: in journalism,

[1]https://ada.silknow.org/
[2]https://explorer.odeuropa.eu/
[3]http://www.kulttuurisampo.fi/?lang=en
[4]https://www.europeana.eu/en

in books, TV-shows and movies, video games, exhibitions, virtual- and augmented-reality environments, etc.[5]

Accessing KG content via stories has been recognised as a complementary approach for knowledge capture and dissemination. This approach has been adopted in several cases such as the DLNarratives project [6], [6] the Belgian Refugees project,[7] the Bletchley Park Text project,[7] the Paths project [8], and the Storyspace project [9].

In this paper, we present an ontology for authoring and presenting hypermedia stories via knowledge graphs, and illustrate how it has been used in practice. The ontology was created in the context of the Odeuropa project, which is focused on the history and heritage of smell [10], but its application generalises to any KG containing object descriptions for which one wants to generate linear and nonlinear stories. We show how existing KG connections can be enriched by story connections, providing contexts that express multiple points of view over a collection of objects in the KG.

The remainder of this paper is organised as follows. In Section 2, we discuss related work in modelling stories. In Section 3 we describe the requirements of the ontology, followed by its implementation in Section 4. We present examples in Section 5, a qualitative evaluation in Section 6 and we conclude with final remarks and directions for future work in Section 7.

The ontology in RDF/Turtle format can be found at https://github.com/Odeuropa/ontology, while its documentation is available at https://data.odeuropa.eu/ontology/story/.

2. Related Work

Prior ontologies for narrative and story structures include StoryTeller [11], BBC News Storyline ontology [12], Labyrinth data model, ODY-ONT [13], Curate ontology [14] and NIAO [15].

Most of these treat stories as linear sequences of events. For example, StoryTeller [11] is an ontology for representing stories as collections of events, which are directly connected with the `contains` property. With this data model, the only way to reconstruct the order of story elements (events) is to compare their time, producing stories which are a pure linear concatenation of episodes, which excludes any descriptive or reflective part of a story. The BBC News Storyline ontology [12] models news topics and events, mostly temporally focused. Similarly, in the Labyrinth data model [16] the *Dynamics* are the events which compose a *Story*. The ODY-ONT model [13] focuses on the representation of events, sub-events, and the temporal relation between them. Similar relations between events are represented in the Datasets of Mystery Stories [17]. The chain of events is also at the foundation of other works [18, 19].

Mulholland et al. [14] present an ontology and system to create unique stories, modeling narratives, plots, stories and events. The *Storyspace API* on top of the *Curate ontology* assists museum curators in manually creating stories, or importing them from Freebase[8].

Automatic narrative generation from data has been explored in various projects. The tourist domain has been explored in [20] in which a mobile guide generates walks along different

---

[5]https://blog.mused.org/an-overview-of-digital-storytelling-platforms-in-2021/

[6]https://dlnarratives.eu/narratives.html

[7]https://tour.belgianrefugees14-18.be/en

[8]Freebase was a large collaborative knowledge base that was shut down in 2016. It now powers part of Google's Knowledge Graph.

cultural points of interest in Turin, Italy. Here, users do not generate the narrative themselves, similarly to [21] in which a timeline of the French revolution is generated automatically from Wikidata. A different view on narratives is presented in [15], which is focused on annotating objects such as murals that describe a multiple-plots story across multiple times and spaces. Here, a narrative consists of one cultural heritage object, instead of multiple, as allowed in our ontology. Extracting storylines from texts has been addressed in the Events and Stories in the news workshops series held from 2015 to 2020 [22] and a handbook compiled by a subset of organisers of this workshop series [23]. Here, events are often expressed as verbs, and tend to be quite fine-grained.

3. Problem Overview

In this section, we give the basic definition and foundation of our work, from which we define the modelling requirements.

3.1. Background concepts

To understand the various layers that make up a storyline, it is useful to refer to concepts based on those proposed by Bal and van Boheemen in [24]. There, a *narrative text* (our storyline) is the production (physical rendition) of a *story*, which in turn is a rhetorical arrangement of elements of a *fabula*. The fabula is a representation of the actual "meaning" of the story, i.e., the conceptualizations that the author intends to create in the reader's mind. A fabula is formed by logically or chronologically related events, and may contain references to external contents. In our ontology, we generalize this notion, considering elements that are not necessarily events.

A *story* can be seen as a collection of atomic *elements* variously referred to in the literature as passages, chunks, fragments, scenes, cards or nodes. A story element has some textual (and, more generally, multimedia) content, and references elements of the underlying KG. It may also reference other sources, e.g., Wikipedia. A story in our definition corresponds to a story in Bal and van Boheemen's terminology, but may have a richer (graph) structure in our case.

In addition, a story establishes structure over its elements, expressed through one or more types of relations between them. These structural relations dictate the order of consumption by the reader. The user starts at an element, and proceeds by consuming (navigating to) any related element according to the story's structure.

From this point of view, the semantics of a KG element is complemented by the semantics of all story elements that reference it, according to the point of view of the story author. In Bal's terminology, the content of the story elements correspond to the fabula.

Once the story is specified, it is necessary to create a presentation, which is the surface form of the story that is effectively consumed by humans. We will discuss this in Section 5.

The most common story topology is a linear structure - the elements form a sequence. This structure is also very common in story representations in the history and cultural heritage domains, specifically chronologically ordered knowledge. However, these structures can be be limited when there is a wish to accommodate more complex narratives, such as the integration of multimedial heritage objects or interactive elements (tests, images, artefacts). Increasingly history and cultural heritage require narratives that incorporate multiple points of view, multiple

time periods, and multiple geographies (multi-perspectivity; pluritemporality). In addition, event-based models struggle to capture more abstract notions (such as complex cultural concepts, collective imaginations, symbolism, or cultural practices) which often form the basis of a historical or cultural heritage storyline.

To illustrate this, we present an example from the history of smell. Consider, for instance the civet cat, a mammal native to the tropical climates of Africa and Asia [25, 26, 27]. At some point, humans discovered and started to appreciate the olfactory qualities of the substance the cat secretes through its perineal gland. The secretion was regarded as having medicinal functions and desirable hedonic qualities. The civet cats 'story' thus changed from a biological story to a story in which the cat also assumed an important role in the history of perfumery, where the role of the cat changed from 'animal' to 'economic product'. From here, 'civet' enters the story of musk, ambergris, and other animal fragrances. Now, the story invites more interest about the human use and evaluation of the civet paste as olfactory substance. With the evolution of social norms and values, its initial hedonic qualities became associated with "lower" social circles such as prostitutes, which also changed its economic value. Civet's use in many perfumes produced from the eighteenth through the early twentieth centuries means that many heritage perfumes being recreated today include civet. Perfumes with a rich heritage, such as Chanel No.5, continue to include a civet note. Meanwhile, a synthetic substitute for its odor element (civetone) was discovered, based on palm oil.

This brief description highlights how one can create a narrative about the civet and civetone from a medicinal perspective (use); from a social perspective (appreciation, symbolism); and from the perspective of abstract concepts such as capitalism. In each perspective (expressed as a storyline), it can be associated with different elements and further storylines, e.g., a story about the economic relevance of palm oil to the perfume industry. Civet's meanings are multi-temporal: the decline in the fashion for civet meant that it obtained an air of pastness, the continuing use of civet in contemporary perfumery ensures its persistence in the present. The attempt to recreate historical perfumes in which civet is used points to its place as a future heritage scent. Smells - and the historical and heritage stories we tell about them - are palimpsests that blend past, present, and future into complex experiences that challenge linear chronologies [28].

From this example, we note that each point of view can be captured in its own linear storyline, with the end result that these various storylines have one or more elements in common. This means that, while consuming the contents of a particular element, the reader will have the option to continue by following along more than one storyline. As a result, the overall story structure effectively becomes a graph.

The next level of complexity in the story structure is when it is natively a graph, regardless of the existence of linear paths through it. This type of structure has been used in hypertext literature[9], games (e.g. Telltale Games[10]) and more recently in interactive film (e.g, the *Black Mirror: Bandersnatch* episode[11]). In these stories, it is left to the reader to choose the available links without a suggested sequence by the author. Here the fabula is only apprehended by the reader after a more extensive exploration of the story elements.

---

[9]https://en.wikipedia.org/wiki/Hypertext_fiction
[10]https://en.wikipedia.org/wiki/Telltale_Games
[11]https://en.wikipedia.org/wiki/Black_Mirror:_Bandersnatch

3.2. Modelling Requirements

In light of the previous discussion, we define the requirements for our proposed ontology:

Req 1: Given the need to accommodate multiple perspectives, and the variety of possible association between elements, it should allow to represent hypertexual (graph-structured) stories, In other words, its elements can be connected through multiple alternative paths;

Req 2: Users can consume stories via navigation links present in each element, in hypertext fashion;

Req 3: To express their particular point-of-view about the elements in a story, authors should be able to define a suggested order of consumption of elements, effectively defining paths (storylines) through the story.

Req 4: An element may be part of different storylines, as the different points-of-view may share common elements. Consequently, while consuming an element, the reader has the option of continuing along the storyline s/he came from or continuing along a different storyline.

Req 5: Stories, while building narratives about different sets of elements, may refer to each other, by sharing common elements - for example, "War" and "Religion" stories. Thus, while consuming one story, the reader should be able to follow a reference to another story.

Req 6: Stories should be able to connect to KGs, and to other sources in general, therefore its elements can refer to external resources via direct links.

Req 7: Since stories reflect points-of-view of its authors, it should be possible to include provenance metadata about the story, namely authors, creation and modification dates, etc.

Req 8: Whereas the semantics of individual concepts referred to in a story element can be given by the underlying KG, there will be additional content in story elements whose semantics are left to the reader to apprehend, as they are expressed through textual and multimedia content. Consequently, the ontology should be independent of (i.e, not make assumptions about) the definition of elements in any underlying KG.

Req 9: A story can be rendered in multiple delivery formats and in multiple environments. It should therefore be neutral in this respect, accommodating this multiplicity.

4. Ontology

We distinguish between the concepts of *Story* and *Storyline*. The basic idea is that the same Story can be told - actually, consumed - in different ways by following different paths throughout the story. Each path (or linearisation, or narrative) is a Storyline which can be chosen in a Story. According to this definition, for each Story there can be multiple Storylines that can be proposed by the storyteller to offer a particular perspective of the Story.

The core part of the StoryLine Ontology is represented in Figure 1. The main building block is the **Story Element**. It is the basic and atomic element of a **Story**. The contents of the element are envisaged to actually be of text or multimedia types, e.g, in HTML, or some structured format to be interpreted by the rendering engine (see Section 5). It is expected that it contains *anchors*, which can be rendered as hypertext links to related (linked) elements which can be navigated by the user. These links are represented in the graph using RDF-Star [29], qualifying (annotating) an RDF statements such as `<element1, connects to, element2>` with the class **Story Connection**. This class has properties for representing its association to a Story and a
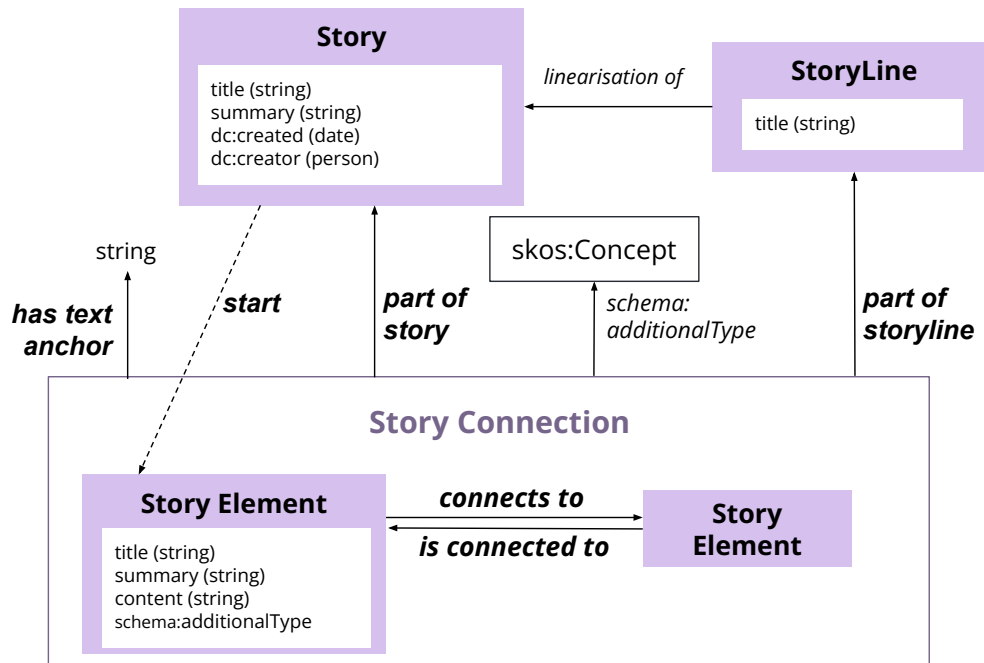
Figure 1: Schema of the StoryLine ontology, with the main elements and the Story Connection

Storyline, and for tracking the text of the anchor in the content to the linked element. These links may have types, which can be used for automatic generation of stories.

These two classes are defined in the following way, keeping in mind that each element can have multiple anchors and, consequently, multiple Story Connections linking to successor elements:

- a *Story* is a collection of interlinked elements, with a starting Story Element and a number of succeeding ones;
- a *StoryLine* identifies a selection of elements that constitute a particular path (set of Story Connection of a Story, one per element) representing one possible *linearisation*.

The presence of both *Story* and *StoryLine* allows to represent both an hypertextual and a linear story, depending on the balance between freedom and guidance that one wants to give to the user. Finally, the ontology allows for an **External Connection** to other entities which may have some relation with the Story Element. We decided to not limit the range of this relation, to enable linking a Story element with knowledge graph entities, but also encyclopaedic articles (e.g., Wikipedia), etc. Also for this connection, it is possible to define an anchor and an additional type, as seen in Figure 2.

We purposefully did not include specific properties besides dc:creator for the metadata of Stories. Given the existence of ontologies that encompass a large number of metadata fields - such as the Dublin Core (dc, as in Figure 1) or the PROV Ontology [30]. We encourage the use of the StoryLine ontology in combination with other schemas, according to the specific needs of the application.
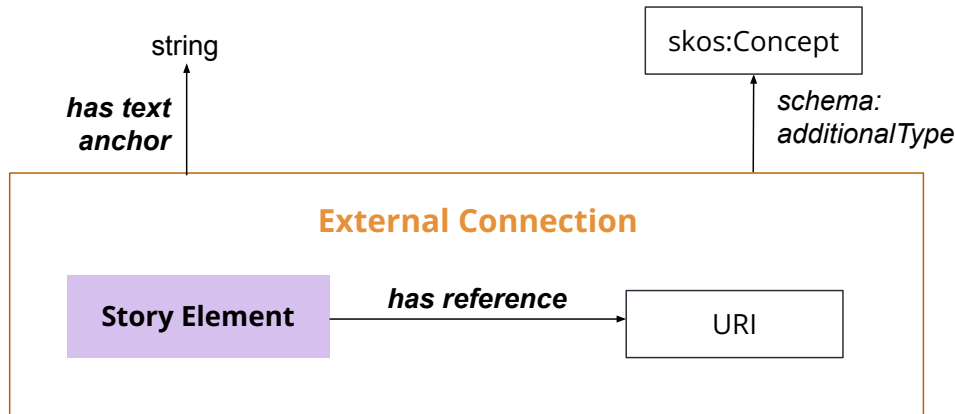
Figure 2: Schema of the StoryLine ontology, with a focus on the External Connection

5. Application Examples

In this section, we present two application examples of our ontology: a hypermedia storyline and a guided city tour application.

Hypermedia storyline

In the Odeuropa project, historians are experimenting with hypermedia narratives to present their syntheses in a data-driven and interactive manner as a complement to traditional monographs. For this, we need to map a story that is represented as an instance of our ontology into a form suitable for human consumption. We refer to this process as "rendering". At the moment, we have taken the approach that the contents of story elements contain presentation specifications that can be directly used by a rendering engine.

In the realm of interactive storytelling tools, Twine[12] is a popular open-source tool which empowers users to create interactive narratives through a visual, node-based interface. In this tool, individual passages or nodes represents story elements, each connected through hyperlinks. These hyperlinks establish the flow of the storyline, allowing readers to make choices that determine the progression of the narrative.

Our Storyline ontology is designed to work well with interactive stories created using tools like Twine, which present stories structured in a graph format – see Figure 3. Moreover, our ontology facilitates bidirectional conversion, enabling the seamless translation of ontology-modeled stories back into Twine-compatible formats.

As an example, a Twine story with the title *War has always been an olfactory affair* is available at https://data.odeuropa.eu/ontology/story/war.html. An excerpt of the RDF version of this story is shown in Listing 1, while a more extended version is available online[13].

---

[12]https://twinery.org/
[13]https://github.com/Odeuropa/ontology/blob/master/storylines_examples/e1.ttl
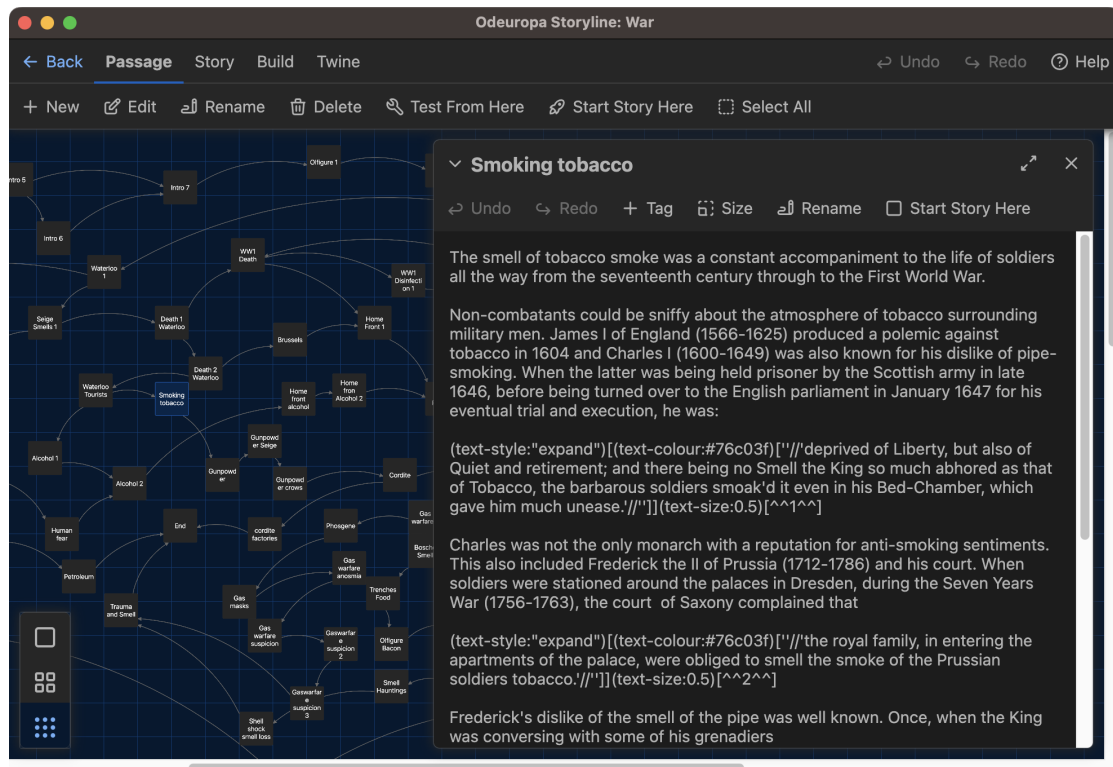
Figure 3: An example of story in Twine, with the story overview on the left and a story element in editing mode on the right. The connections between cards (the story elements) naturally give to the story a graph structure.

Listing 1: The representation in RDF of the War story

```
:War a story:Story;
        story:title "War has always been an olfactory affair";
        story:start :Introduction .
:War1 a story:Story;
        story:linearisationOf :War ;
        story:title "War (short version)" .
:waterloo_1 a story:StoryElement;
    story:title "Waterloo";
        story:content """ ... """.
:waterloo_1 story:connectsTo :seige_smells_1 {|
        a story:StoryConnection ;
        story:hasTextAnchor "even more troubling smells" ;
        story:partOfStory :War ;
        story:partOfStoryline :War1 |}.
:waterloo_1 story:connectsTo :incense {|
        a story:StoryConnection ;
        story:hasTextAnchor "Incense" ;
        story:partOfStory :Religion . }.
```

(a) 'Rub'n'Sniff' map created for Odeuropa's, *City Sniffers-A smell tour of Amsterdam's ecohistory*, printed for the event by Scent the Brand. Designed by Liam R. Findlay. Note: the numbers indicate the order of stops.

(b) Screen shot of the City Sniffers application (map).

(c) Screen shot of the City Sniffers application (stop 2).

Figure 4: City Sniffers 'Rub and Sniff' card and smartphone application screenshots.

City Sniffers Guided Tour

The City Sniffers application[14] is a smell tour of Amsterdam's eco-history. It consists of a rub and sniff card and a mobile phone application. Users can follow a path of six selected stops through Amsterdam, the Netherlands. At each stop, the participant scratches the corresponding spot on the card and can sniff a scent while reading a text explaining the smell and its cultural-historical context and relevance in the application. Figure 4 shows the rub and sniff card alongside two screenshots of the applications. This application also links to the Odeuropa KG and story elements are the texts associated with each stop. The storyline structure is the linear path connecting each spot to its successor.

6. Evaluation

We present a qualitative evaluation of the ontology, by examining – together with researchers from the humanities domain – each of the requirements given in Section 3.2 and showing how they are satisfied. To facilitate reading, each item mentions its corresponding requirement.

Req 1: The graph-like structure is enabled by the "connects-to" relation, and allows an element to participate in several relation instances.
Req 2: Hypertext navigation by the user is enabled by the "anchor" property, which is used by the rendering environment to allow navigation;

---

[14] https://citysniffers.odeuropa.eu/

Req 3: Storylines are enabled via the association of a link to the "StoryLine" class in the ontology.

Req 4: An element may be part of different storylines, since "connects-to" relation can be qualified by more than one "Storyline" instance in the ontology.

Req 5: Stories may refer to each other since the source and the destination elements in a "connects-to" relation may each also be part of a different "connects-to" relation that are respectively qualified with different instances of "Story".

Req 6: Stories are be able to connect to KGs, and to other sources, since the "Story Element" class has a "has reference" property that points an external URL.

Req 7: Stories can have provenance meta-data, although only the "dc:creator" property of the "Story" class is explicitly included in the ontology.

Req 8: The ontology is independent of the definition of elements in any underlying KG since the content of the "Story Element" class has a "content" property which is a string. This string can be interpreted in innumerable ways by the rendering environment.

Req 9: A story can be rendered in multiple formats and multiple environments, since the "content" and "anchor" attributes may contain formatting instructions to be interpreted by the rendering environment, such as Twine (which uses a markup language) .

7. Conclusion and Future Work

In this paper, we have presented an ontology for creating hypermedia stories over knowledge graphs using RDF-Star and shown its use in two example applications. Our ontology effectively enables a narrative layer over large datasets. Our ontology allows smooth integration with underlying KGs, and enables stories with rich hypermedia structure, encompassing non-event-based elements as well.

In future work, we aim to further expand the ontology to include provisions for expressing rhetorical structures such as the ones enabled by the Curate ontology [14], and others. Furthermore, we will refine the provenance requirements to aid users in tracing their information sources and narrative development. Additionally, we plan to realise an automatic converter between the Twine format and the Storyline ontology representation, and vice versa to empower humanities and GLAM professionals to create storylines to expose their archival searches and collections. We also plan to extend the implementation to allow users to record, resume and possibly share navigation histories over a set of stories.

Acknowledgments

References

[1] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula,

L. Schmelzeisen, J. F. Sequeda, S. Staab, A. Zimmermann, Knowledge graphs, ACM Comput. Surv. 54 (2022) 71:1–71:37. doi:10.1145/3447772.

[2] T. Ehrhart, P. Lisena, R. Troncy, KG explorer: a customisable exploration tool for knowledge graphs, in: VOILA 2021, International Workshop on the Visualization and Interaction for Ontologies and Linked Data, 2021.

[3] E. Mäkelä, E. Hyvönen, T. Ruotsalo, How to deal with massively heterogeneous cultural heritage data–lessons learned in CultureSampo, Semantic Web 3 (2012) 85–109.

[4] V. Petras, T. Hill, J. Stiller, M. Gäde, Europeana–a search engine for digitised cultural heritage material, Datenbank-Spektrum 17 (2017) 41–46.

[5] K. Dalkir, E. Wiseman, Organizational Storytelling and Knowledge Management: A Survey, Storytelling, Self, Society 1 (2004) 57–73. URL: https://www.tandfonline.com/doi/abs/10.1080/15505340409490258.

[6] C. Meghini, V. Bartalesi, D. Metilli, Representing Narratives in Digital Libraries: The Narrative Ontology, Semant. Web 12 (2021) 241–264. doi:10.3233/SW-200421.

[7] P. Mulholland, T. Collins, Z. Zdráhal, Bletchley Park Text: Using mobile and semantic web technologies to support the post-visit use of online museum resources, Journal of Interactive Media in Education 2005 (2005). doi:10.5334/2005-24.

[8] K. Fernie, J. Griffiths, M. Stevenson, P. Clough, P. Goodale, M. Hall, P. Archer, K. Chandrinos, E. Agirre, O. de Lacalle, A. de Polo, R. Bergheim, PATHS: Personalising access to cultural heritage spaces, in: 2012 18th International Conference on Virtual Systems and Multimedia, 2012, pp. 469–474. doi:10.1109/VSMM.2012.6365960.

[9] A. Wolff, P. Mulholland, T. Collins, Storyspace: A Story-Driven Approach for Creating Museum Narratives, in: Proceedings of the 23rd ACM Conference on Hypertext and Social Media, HT '12, Association for Computing Machinery, New York, NY, USA, 2012, p. 89–98. doi:10.1145/2309996.2310012.

[10] P. Lisena, M. van Erp, C. Bembibre, I. Leemans, Data Mining and Knowledge Graphs as a Backbone for Advanced Olfactory Experiences, in: STT21: Smell, Taste, and Temperature Interfaces workshop, Yokohama, Japan, 2021. URL: https://stt21.plopes.org/wp-content/uploads/2021/05/STT2021_Odeuropa.pdf.

[11] J.-h. Yeh, StoryTeller: An event-based story ontology composition system for biographical history, in: 2017 International Conference on Applied System Innovation (ICASI), 2017, pp. 1934–1937. doi:10.1109/ICASI.2017.7988448.

[12] P. Wilton, J. Tarling, J. McGinnis, Storyline Ontology - An ontology to represent News Storylines., 2013. URL: https://www.bbc.co.uk/ontologies/storyline-ontology/.

[13] A. F. Khan, A. Bellandi, G. Benotto, F. Frontini, E. Giovannetti, M. Reboul, Leveraging a Narrative Ontology to Query a Literary Text, in: B. Miller, A. Lieto, R. Ronfard, S. G. Ware, M. A. Finlayson (Eds.), 7th Workshop on Computational Models of Narrative (CMN 2016), volume 53 of *OpenAccess Series in Informatics (OASIcs)*, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016, pp. 10:1–10:10.

[14] P. Mulholland, A. Wolff, T. Collins, Curate and Storyspace: An Ontology and Web-Based Environment for Describing Curatorial Narratives, in: E. Simperl, P. Cimiano, A. Polleres, O. Corcho, V. Presutti (Eds.), The Semantic Web: Research and Applications, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 748–762.

[15] L. Xu, A. Merono-Penuela, Z. Huang, F. Van Harmelen, An ontology model for narrative

image annotation in the field of cultural heritage, in: Proceedings of the 2nd Workshop on Humanities in the Semantic web (WHiSe), 2017, pp. 15–26.

[16] R. Damiano, A. Lieto, Ontological Representations of Narratives: a Case Study on Stories and Actions, in: M. A. Finlayson, B. Fisseni, B. Löwe, J. C. Meister (Eds.), 2013 Workshop on Computational Models of Narrative, volume 32 of *OpenAccess Series in Informatics (OASIcs)*, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2013, pp. 76–93. doi:`10.4230/OASIcs.CMN.2013.76`.

[17] K. Kozaki, S. Egami, K. Matsushita, T. Ugai, T. Kawamura, K. Fukuda, Datasets of Mystery Stories for Knowledge Graph Reasoning Challenge, in: SEMMES'23: Semantic Methods for Events and Stories (ESWC2023), Hersonissos, Greece, 2023.

[18] W. Croft, P. Pešková, M. Regan, Integrating Decompositional Event Structures into Storylines, in: Proceedings of the Events and Stories in the News Workshop, Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 98–109. URL: https://aclanthology.org/W17-2713. doi:`10.18653/v1/W17-2713`.

[19] J. Moreno-Schneider, P. Bourgonje, G. Rehm, Towards User Interfaces for Semantic Storytelling, in: S. Yamamoto (Ed.), Human Interface and the Management of Information: Supporting Learning, Decision-Making and Collaboration, Springer International Publishing, Cham, 2017, pp. 403–421.

[20] N. Mauro, A. Geninatti Cossatin, E. Cravero, L. Ardissono, G. Magnano, M. Giardino, Exploring Semantically Interlaced Cultural Heritage Narratives, in: Proceedings of the 33rd ACM Conference on Hypertext and Social Media, 2022, pp. 192–197.

[21] I. Blin, Building a French Revolution Narrative from Wikidata, in: JCAI 2022: Workshop on semantic techniques for narrative-based understanding, 2022.

[22] C. Bonial, T. Caselli, S. Chaturvedi, E. Clark, R. Huang, M. Iyyer, A. Jaimes, H. Ji, L. J. Martin, B. Miller, T. Mitamura, N. Peng, J. Tetreault (Eds.), Proceedings of the First Joint Workshop on Narrative Understanding, Storylines, and Events, Association for Computational Linguistics, Online, 2020. URL: https://aclanthology.org/2020.nuse-1.0.

[23] T. Caselli, E. Hovy, M. Palmer, P. Vossen, Computational Analysis of Storylines: Making Sense of Events, Cambridge University Press, 2021.

[24] M. Bal, C. van Boheemen, Narratology: Introduction to the Theory of Narrative - 4th. Ed., University of Toronto Press, 2017.

[25] K. H. Dannenfeldt, Europe discovers civet cats and civet, Journal of the History of Biology (1985) 403–431.

[26] H. Dugan, The ephemeral history of perfume: Scent and sense in early modern England, JHU Press, 2011.

[27] E. Fudge, Renaissance animal things, New Formations 76 (2012) 86–100.

[28] W. Tullett, Smell and the Past: Noses, Archives, Narratives, Bloomsbury, 2023.

[29] O. Hartig, P.-A. Champin, G. Kellogg, A. Seaborne, RDF-star and SPARQL-star: Final Community Group Report, https://w3c.github.io/rdf-star/cg-spec/2021-12-17.html, 2021.

[30] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, PROV-O: The PROV ontology, Technical Report, World Wide Web Consortium, 2013.