On Design Principles for Narrative Information Systems

Hermann Kroll[1], Wolf-Tilo Balke[1]

[1]*Institute for Information Systems, TU Braunschweig, Mühlenpfordtstr. 23, Braunschweig, 38106, Germany*

Abstract

Information systems enable users to represent, manipulate, and access data effectively. Currently, the development and availability of extensive knowledge bases receive a lot of attention for supporting a variety of *intelligent* applications like smart assistants, question answering, and knowledge discovery. But building such knowledge bases generally means breaking down information into manageable pieces (e.g., in the form of triples in RDF) that tend to lose connections and extraction contexts. In contrast, humans typically share and exchange knowledge by interweaving different pieces into *narratives* that are *plausible* and *easy to grasp*. In this position paper, we summarize the main research directions of narratives in computer sciences. Moreover, we propose basic design principles that narrative information systems should consider in practice. In particular, we take a closer look at *narrative representations*, possible *bindings* between narratives and real-world data, the *context-compatibility* of information, and finally a narrative's *plausibility*.

Keywords

Narratives, Information Systems, Plausibility, Context, Knowledge Representation

1. Introduction to Narratives

Supporting *narratives* for knowledge exchange is a new challenge for information systems research, as humans tend to process information easier whenever they are presented as (a part of) coherent narratives [1]. The term *narrative* is, however, hard to grasp: Some understand the term narrative synonymously to the concept of a story, in essence referring to a temporally or causally ordered sequence of events. Some specifically distinguish between fictional and non-fictional narratives. Others talk about the intention of a narrative, e.g., to convince others of a particular stance.

Our research is focused on the exchange of knowledge through meaningful patterns, i.e., we understand a narrative as a possible form to share knowledge. In this paper, we focus on the exchange and transfer of knowledge through plausible narratives. Therefore, we first examine how established systems represent knowledge and discuss their limitations. We continue with a brief introduction to narratives in computer science and move forward to our position: Narratives need to be handled as first-class citizens in information systems.

Knowledge Bases. So how do current systems represent knowledge? In contrast to classical relational databases or warehouses for well-structured data, the development of the Resource Description Framework (RDF) within the standardization efforts of the Semantic Web has allowed a structured and simple-to-use repre-

sentation of (mostly entity-centric) knowledge [2] that is adapted in most of today's knowledge bases. The central idea is to store knowledge in the form of triples, i.e., subject-predicate-object tuples like *(Albert Einstein, was born in, Germany)*. These triples are often called *facts* or *statements*. A collection of such facts is then called a knowledge base. In recent years, extensive knowledge bases have been collected mostly from Web data: The Wikidata [3] project, for example, includes knowledge from many different domains. Wikidata users can either access the knowledge in a browsing fashion or retrieve it using the structured query language SPARQL.

Besides storing knowledge about entities, there are also knowledge bases that capture knowledge about other types of things. In particular, there is an increasing interest in knowledge bases focusing on events, e.g., Knowlywood captures information about activities mined from Hollywood narratives [4], and the EventKG captures all kinds of information about a variety of events [5]. To efficiently build such collections, methods have been proposed to extract events from textual sources, e.g., events from news [6], temporal facts and events [7], or events from Wikipedia [8]. For a good overview on event extraction, see [9].

In brief, the resulting knowledge bases store factual knowledge about entities or events, e.g., an entity's name, age, and type or some event's date, location, and participants. Moreover, it also reflects typed relations between pairs of entities and events. So is this knowledge representation sufficient for information systems? From our perspective, knowledge bases indeed effectively represent and allow access to *factual knowledge*, in the sense of universally valid statements. In strict correspondence to the real world, such statements can consistently co-exist within a knowledge base. But what happens to knowledge that is *not universally valid*? For instance,
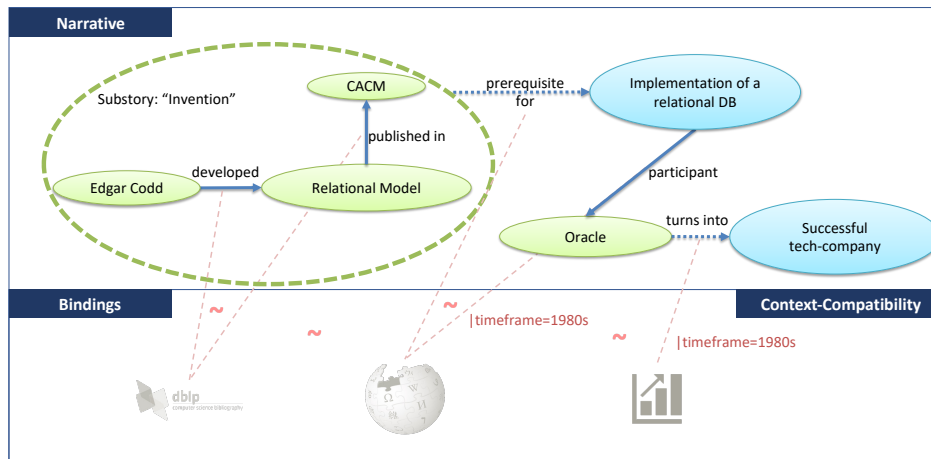
Figure 1: Design Principles for Narrative Information Systems by an example: The narrative structures the story that Oracle has turned into a successful tech-company after implementing a relational database, for which the development and publishing of the relational model by Edgar Codd was an essential prerequisite. The different relationships can be bound against knowledge of DBLP, Wikipedia, and stock values. The bindings share a common context in the sense of time, i.e., *timeframe 1980s* (the remaining bindings are valid in general and thus are compatible too). These bindings are thus context-compatible and cover the whole narrative – making it plausible.

some statements may either reflect subjective opinions or they may be only valid in certain contexts or under certain conditions or even only with a certain probability. The resulting knowledge containing also such facts base might easily become inconsistent or contradictory. Thus, while knowledge bases are effective for factual knowledge, they suffer in representing the connection between statements.

A good overview of the limitation of today's knowledge bases is given in [10]. Suchanek argues for moving beyond triple-based structures to support negations, dynamic behavior, and beliefs. For example, a certain statement might be accepted by a group of people and rejected by some other group: how to deal with both pieces of information within the same repository? As another example take two individual statements being prerequisite and reason for a third statement: currently this cannot be represented in a knowledge base, at least not with a triple-like structure. Suchanek argues for representing such situations as *frames*, i.e., scopes that connect knowledge pieces in different ways. In contrast, we argue to reconsider on how humans exchange knowledge: through plausible narratives.

Narratives. Narratives have been studied in computer science for a while. Sequences of events or actions to model human behavior and situation awareness go back to Schank and Abelson [11] in the form of scripts. [12] proposed the representation of narratives as a *sequence of events*: Narrative schemas have been introduced to represent a set of events and their semantic roles [13]

and were collected in a respective schema database [14].

Previous works generally assume an order and relation of events in texts that needs to be mined. Tang et al. proposed neural methods to learn the respective relations between events from text directly [15]. Another approach proposed the mining of so-called plot graphs [16]. In contrast to describing plots as conjunctions of events, such plot graphs also allowed disjunctions between events, optional, and mutually exclusive events. In brief, a plot graph is a template to craft a certain narrative, i.e., the plot graph determines which sequences of events are allowed.

Narrative Applications. Various applications have been proposed to utilize narratives in information systems. Narrative schemas and plot graphs are used for *narrative planning*, i.e., to write specific stories by computers [17, 18]. Another application is the alignment of narratives and movie scripts, i.e., linking narrative components (e.g., *fall in love*) to performed actions (e.g., *A kisses B*) in movies [19].

The digital library project Europeana utilizes narratives to compose events involving artists into timelines [20]. The narratives are then shown to users as overviews of certain artists, i.e., to tell a story about them. The benefits of such a representation have already been studied [21]. Moreover, knowledge bases like Wikidata can be utilized to populate narratives with additional factual information [22].

Finally, Shahaf et al. propose to connect the dots between news articles in so-called *metro maps* [23]. Metro

maps consist of articles as stations and topics as metro lines. Users may then follow a line (topic) across different articles. In addition to news items, Shahaf et al. have demonstrated how scientific articles can be arranged as metro maps of sciences, too [24]. In summary, such Metro maps allow users to understand better the spatial, causal, or temporal structure of available resources. We understand such maps as a possible application of narratives aligning similar textual sources in a meaningful way.

2. Narratives and Plausibility

Our main objective is to integrate narratives in information systems to allow a *meaningful* representation of knowledge. Therefore, we discuss four central design principles that narrative information systems must consider: narrative representation, bindings, context-compatibility, and plausibility.

When is a narrative *meaningful*? While we agree that *meaningfulness* is strongly connected to the actual structure of a narrative, our research is focused on the *plausibility* of narratives. We argue that plausible narratives have a *meaning*. Consider our example in Fig. 1: The development of the relational model by Edgar Codd was a prerequisite for the implementation of a relational database system by Oracle. Oracle has then turned into a successful tech-company. The narrative is considered plausible because of two reasons: First, all of its parts have evidence in knowledge repositories. Second, the *timeframe* of the implementation matches the *timeframe* of Oracle's stock value increase. With that in mind, we first argue on narrative representations and then continue with key concepts to assess the plausibility of a narrative.

Narrative Representation. We understand a narrative as a structure that describes some progress, e.g., of temporal nature (one event might follow another event), of conditional nature (one event might depend on the other event being observed before), or of causal nature (one event leads to another event). In previous work, we designed narratives as logical overlays on top of knowledge repositories [25]. Here, we summarize what a narrative representation should consider.

Components of a narrative are **entities** (relevant things of interest, concepts, etc.), **events** (something which happens, some observed state, a change of some state, etc.) and **literals** (values). We made the distinction between entities and events to distinguish between rather *static* things from the real world and *dynamic* happenings that have a temporal dimension or represent some state or state change.

A narrative then puts these components in relation to each other. Therefore, we distinguish between two different kinds of relationships: **factual** and **narrative relationships**. Like knowledge bases, factual relationships describe properties involving entities and literals, e.g., that *Edgar Codd developed the relational model*. These relationships can be put between entities and literals. We compose the actual progress of a narrative by a set of narrative relationships, e.g., *the implementation of a relational database system* by *Oracle* that has since *turned into* a *successful tech-company*. In brief, narrative relations feature special, non-factual labels. They can be placed between events or between events and entities, but not between entities. In this way, narrative can represent static knowledge and interweave such facts into dynamic progress (e.g., some causal structure, some temporal flow, etc.).

The last thing a narrative representation must consider is *nesting*. With nesting, also called recursion or induction, we refer to using narratives inside narratives, e.g., a single narrative leads to another narrative. Nesting supports recursive elements as components, known from spoken languages [26]. In our example, the *development of the relational model* forms a narrative on its own that is, in turn, the prerequisite for the actual implementation. While we agree that nesting makes the representation more challenging, especially concerning precise semantics, we argue that nesting allows a higher expressiveness of narratives, e.g., the occurrence of narrative a's events together finally lead to the whole narrative b.

Narrative Bindings. For now, narratives are *rather* artificial structures. However, the question remains, how can we estimate whether the narrative is *plausible*? We argue that evidence for a narrative is one way to go: Evidence allows us to estimate if a narrative is *grounded* by some real-world data. If a structure is based on real-world data, the narrative itself is supported and has thus evidence. Such a *grounding* leads to a better understanding of whether the narrative is *plausible*.

The next step for a *plausibility assessment* is thus to connect narratives with knowledge repositories. In our understanding, a knowledge repository can be any collection of knowledge, be it a collection of articles, a structured database, a data set repository, etc. We call the connection between a narrative's relationship and a knowledge repository a **binding** [25]. In other words, the binding binds a certain relationship against real-world data. It gives, in this way, evidence for it. In our example, the development of the relational model by Edgar Codd can be bound against DBLP (a bibliographic database of computer science). The implementation of a relational database by Oracle could be bound against a Wikipedia article. In brief, different knowledge repositories may require different computations of bindings [27]. To give a few examples:

Structured Repositories like databases and knowledge bases usually support structured query languages such as SPARQL, Gremlin, or SQL: Bindings for those knowledge sources may be computed by translating narrative components to respective queries.

Relational Data Sets Tabular data formats comprise a large amount of knowledge. However, the computation of bindings against those data requires a certain amount of knowledge about metadata [28]. A binding must first determine which columns of a table represent a component in a narrative and afterward find a statistical measurement to determine if a narrative's relationship is supported. In our example, an application of this procedure is the measurement of the success of Oracle against stock values.

Textual Sources Computing bindings for textual sources can roughly be done in two ways: 1. preprocessing of the text into a structured representation (e.g., graphs) [29] or 2. text-based retrieval methods like traditional keyword-based retrieval or textual entailment methods. Especially domains that can be described by controlled vocabularies (e.g., the biomedical domain), benefit from graph representations [29].

Context-Compatibility. Briefly speaking, bindings allow us to find evidence for a narrative. But can we just compute bindings for each relationship separately to ground the whole narrative? While some pieces of information are universally applicable and easy to connect to other pieces, such as the *birth date* of some person, there may be pieces that are only true within specific semantic settings. For instance, the *capital* of some country may change over time and thus should only be connected to pieces valid in the same time frame. Thus, when extracting pieces of information from knowledge repositories, it is necessary to consider potential **contexts**. Contexts are given by constraints on environment variables that describe under which condition a certain piece of information is valid.[1]

Narrative information systems must consider contexts to support a valid information fusion of pieces, i.e., to fuse only pieces that are valid under the same context conditions. In our example, the turning into a successful tech-company should be valid in the same timeframe as the implementation of the actual relational database. We call bindings context-compatible if they bind relationships of narratives against knowledge that is valid under the same context conditions.

The question remains how contexts can be represented: On the one hand, explicit context models allow one to model every single condition explicitly. For example, McCarthy proposed a model based on first-order predicate logic [30]. Temporal-restricted or spatial-restricted statements might be a good example of such explicit models. Suppose each statement is attached with a temporal interval defining the statement's validity. In that case, a reasoning process could then only consider statements that are valid in the same time interval, e.g., a compatible timeframe as introduced in our running example.

On the other hand, explicitly modeling every condition can be cost-intensive and, in some domains, may be close to impossible. We, therefore, proposed an implicit context model [31]: If statements are used in textual sources, e.g., scientific publications or articles, we assume that these articles should implicitly include the relevant context conditions. We then proposed textual or metadata-based (e.g., authors) metrics to estimate whether two statements are context-compatible, e.g., if the articles, they belong to, are similar enough.

While explicit models allow a controlled fusion, in the sense of correctness, they might be cost-intensive to develop and maintain. Implicit models are cheaper and easier to use. However, they do not guarantee correctness and may lag behind in quality and explainability. For an extensive discussion on contexts, we refer the reader to our previous article on context models [32].

Plausibility. With the basic requirements of narrative structures, bindings, and context-compatibility discussed, we can now define when we call a narrative plausible.

Definition 1. *We call a narrative n plausible, iff the following conditions hold:*

1. *There exists a set of bindings NB that bind each relationship of n against real-world data.*
2. *This set of bindings NB must be context-compatible, i.e., there is at least a single context to which all bindings agree (are valid in).*

We call the task to find such an answer narrative query processing [32]. The central idea is that if we can bind the whole narrative structure, and all bindings are also context-compatible, then there is decisive evidence that the narrative is plausible – again in the sense of having evidence.

However, although we defined narrative plausibility, there still remain challenges for a practical application. We proposed four major dimensions that influence the overall plausibility of a narrative [33]:

---

[1]Please note that for the scope of this paper, we are abstracting from the actual *truth* or *correctness* of each piece of information, which may rather reflect on the respective repository's trustworthiness than provide a context in the sense of validity.

Narrative Structure: A slight change of a narrative, e.g., adding, editing, or leaving out some component, can affect the result of a plausibility assessment, i.e., different narratives might end up in different assessments (plausible/not plausible).

Validation Approach: Which knowledge repositories, and especially, which methods are used to compute the bindings? What do the methods guarantee? And thus, what could a user then expect?

Types of Evidence: Are we just looking for direct bindings? What about counter-examples, i.e., statements that contradict a narrative's relationship?

Confidence of Bindings: Each binding should be connected to confidence, i.e., a score that describes how certain we are that the relationship can be bound against some knowledge repository. Therefore, confidence is related to two properties: the trustworthiness of knowledge repositories and the quality of the binding process (e.g., the confidence of a retrieval or NLP method).

3. Narrative Information Systems

We do not believe there will be a unified narrative model that rules all purposes. Nevertheless, we have seen decisive properties that should be considered. From our perspective, a **narrative information system** must thus consider:

Narrative Representation that makes a distinction between factual knowledge (e.g., entities and their properties) and dynamic progress (e.g., causation and temporality).

Bindings that connect artificial narrative structures with real-world data and events to give evidence for the narrative.

Context-Compatibility that embeds information pieces into contexts and defines rules when contexts are compatible to support a valid information fusion.

Plausibility of narratives by binding the whole narrative with context-compatible bindings.

So, why should we implement a narrative information system? In contrast to pure knowledge bases, narrative information systems allow us to interweave factual knowledge into plausible patterns. This way, the knowledge is represented in a pattern like humans would exchange their thoughts. Reaching this target then allows a set of useful applications.

Narrative Information Access. Narrative information access allows user to formulate their information need as a narrative. Then query processing takes place to make the narrative plausible, i.e., finding context-compatible bindings that cover the whole narrative. We already demonstrated the benefits of narrative information access to digital libraries in [32, 29].

Linking Data Sets to Narratives. Today, research data management has become more apparent than ever. But connecting published data sets and claims from scientific publications remains challenging. If we could link both, we, on the one hand, could explain what a data set *tells*, and on the other hand, verify claims of new publications by existing data sets [28].

Narrative Event Aspects. For now, the previous argumentation focused on the exchange of knowledge through narratives. How a narrative, or its events, are *perceived* by an audience also plays a central role today (think about fake news and social media perception here). Attributions and roles could further extend the narrative representation for richer semantics to better describe these events [34, 35].

4. Conclusions

The design and implementation of narrative information systems promise to support a wide range of novel and exciting applications to support human-centered workflows, e.g., by accessing and explaining knowledge through plausible narratives in the sense of storytelling. Knowledge shared in this way is bound to allow for a better and easier understanding of how a piece of knowledge is placed in the respective domain. And of course, such models could even be extended further, for instance, by adding attributions also to support subjective aspects of knowledge in the sense of opinions, reflect special user intents in how knowledge can or should be used, or express emotional stances towards entities, events, and temporal or causal developments.

In summary, the central argument of this paper is that in order to build effective narrative information systems, we need to treat narratives as first-class citizens within these systems. Arguably there might not be a single model that applies perfectly to all use cases. Nevertheless, there are some key properties and essential problems that every narrative information system needs to consider, such as extraction methods, representational issues, or the possibility and validity of information fusion. This paper opens up the relevant design dimensions and points to some early-stage solutions.

5. Acknowledgments

References

[1] J. László, The Science of Stories: An Introduction to Narrative Psychology, Routledge, 2008. doi:10.4324/9780203894934.

[2] F. Manola, E. Miller, B. McBride, et al., Rdf primer, W3C recommendation 10 (2004) 6.

[3] D. Vrandecic, M. Krötzsch, Wikidata: a free collaborative knowledgebase, Commun. ACM 57 (2014) 78–85. doi:10.1145/2629489.

[4] N. Tandon, G. de Melo, A. De, G. Weikum, Knowlywood: Mining Activity Knowledge From Hollywood Narratives, in: Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM '15, Association for Computing Machinery, New York, NY, USA, 2015, p. 223–232. doi:10.1145/2806416.2806583.

[5] S. Gottschalk, E. Demidova, EventKG - the hub of event knowledge on the web - and biographical timeline generation, Semantic Web 10 (2019) 1039–1070. doi:10.3233/SW-190355.

[6] E. Kuzey, J. Vreeken, G. Weikum, A Fresh Look on Knowledge Bases: Distilling Named Events from News, in: Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM '14, Association for Computing Machinery, New York, NY, USA, 2014, p. 1689–1698. doi:10.1145/2661829.2661984.

[7] E. Kuzey, G. Weikum, Extraction of Temporal Facts and Events from Wikipedia, in: Proceedings of the 2nd Temporal Web Analytics Workshop, TempWeb '12, Association for Computing Machinery, New York, NY, USA, 2012, p. 25–32. doi:10.1145/2169095.2169101.

[8] A. Spitz, M. Gertz, Terms over LOAD: Leveraging Named Entities for Cross-Document Extraction and Summarization of Events, in: Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016, ACM, 2016, pp. 503–512. doi:10.1145/2911451.2911529.

[9] W. Xiang, B. Wang, A Survey of Event Extraction From Text, IEEE Access 7 (2019) 173111–173137. doi:10.1109/ACCESS.2019.2956831.

[10] F. M. Suchanek, The Need to Move beyond Triples, in: Proceedings of Text2Story - Third Workshop on Narrative Extraction From Texts co-located with 42nd European Conference on Information Retrieval, Text2Story@ECIR 2020, Lisbon, Portugal, April 14th, 2020 [online only], volume 2593 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020, pp. 95–104. URL: http://ceur-ws.org/Vol-2593/paper12.pdf.

[11] R. Schank, R. Abelson, Scripts, Plans, Goals, and Understanding, Routledge, 1977.

[12] N. Chambers, D. Jurafsky, Unsupervised Learning of Narrative Event Chains, in: ACL 2008, Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics, June 15-20, 2008, Columbus, Ohio, USA, The Association for Computer Linguistics, 2008, pp. 789–797. URL: https://aclanthology.org/P08-1090/.

[13] N. Chambers, D. Jurafsky, Unsupervised Learning of Narrative Schemas and their Participants, in: ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2-7 August 2009, Singapore, The Association for Computer Linguistics, 2009, pp. 602–610. URL: https://aclanthology.org/P09-1068/.

[14] N. Chambers, D. Jurafsky, A Database of Narrative Schemas, in: Proceedings of the International Conference on Language Resources and Evaluation, LREC 2010, 17-23 May 2010, Valletta, Malta, European Language Resources Association, 2010. URL: http://www.lrec-conf.org/proceedings/lrec2010/summaries/58.html.

[15] J. Tang, H. Lin, M. Liao, Y. Lu, X. Han, L. Sun, W. Xie, J. Xu, From Discourse to Narrative: Knowledge Projection for Event Relation Extraction, in: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Association for Computational Linguistics, Online, 2021, pp. 732–742. doi:10.18653/v1/2021.acl-long.60.

[16] B. Li, S. Lee-Urban, G. Johnston, M. O. Riedl, Story Generation with Crowdsourced Plot Graphs, in: Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, July 14-18, 2013, Bellevue, Washington, USA, AAAI Press, 2013. URL: http://www.aaai.org/ocs/index.php/AAAI/AAAI13/paper/view/6399.

[17] J. R. Meehan, The Metanovel: Writing Stories by Computers, Yale University, 1976.

[18] Mark O. Riedl and Robert Michael Young, Narrative Planning: Balancing Plot and Character, J. Artif. Intell. Res. 39 (2010) 217–268. doi:10.1613/jair.2989.

[19] P. Mirza, M. Abouhamra, G. Weikum, AlignNarr: Aligning Narratives on Movies, in: Proceedings

of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 2: Short Papers), Virtual Event, August 1-6, 2021, Association for Computational Linguistics, 2021, pp. 427–433. doi:10.18653/v1/2021.acl-short.54.

[20] C. Meghini, V. Bartalesi, D. Metilli, F. Benedetti, Introducing Narratives in Europeana: Preliminary Steps, in: New Trends in Databases and Information Systems - ADBIS 2017 Short Papers and Workshops, AMSD, BigNovelTI, DAS, SW4CH, DC, Nicosia, Cyprus, September 24-27, 2017, Proceedings, volume 767 of *Communications in Computer and Information Science*, Springer, 2017, pp. 333–342. doi:10.1007/978-3-319-67162-8_33.

[21] C. Meghini, V. Bartalesi, D. Metilli, F. Benedetti, Introducing narratives in Europeana: A case study, Int. J. Appl. Math. Comput. Sci. 29 (2019) 7–16. doi:10.2478/amcs-2019-0001.

[22] D. Metilli, V. Bartalesi, C. Meghini, A Wikidata-based tool for building and visualising narratives, Int. J. Digit. Libr. 20 (2019) 417–432. doi:10.1007/s00799-019-00266-3.

[23] D. Shahaf, C. Guestrin, Connecting the dots between news articles, in: Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, July 25-28, 2010, ACM, 2010, pp. 623–632. doi:10.1145/1835804.1835884.

[24] D. Shahaf, C. Guestrin, E. Horvitz, Metro maps of science, in: The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, Beijing, China, August 12-16, 2012, ACM, 2012, pp. 1122–1130. doi:10.1145/2339530.2339706.

[25] H. Kroll, D. Nagel, W. Balke, Modeling Narrative Structures in Logical Overlays on Top of Knowledge Repositories, in: Conceptual Modeling - 39th International Conference, ER 2020, Vienna, Austria, November 3-6, 2020, Proceedings, volume 12400 of *Lecture Notes in Computer Science*, Springer, 2020, pp. 250–260. doi:10.1007/978-3-030-62522-1_18.

[26] M. D. Hauser, N. Chomsky, W. T. Fitch, The Faculty of Language: What Is It, Who Has It, and How Did It Evolve?, Science 298 (2002). doi:10.1126/science.298.5598.1569.

[27] H. Kroll, D. Nagel, M. Kunz, W. Balke, Demonstrating Narrative Bindings: Linking Discourses to Knowledge Repositories, in: Proceedings of Text2Story - Fourth Workshop on Narrative Extraction From Texts held in conjunction with the 43rd European Conference on Information Retrieval (ECIR 2021), Lucca, Italy, April 1, 2021 (online event due to Covid-19 outbreak), volume 2860 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2021, pp. 57–63. URL: http://ceur-ws.org/Vol-2860/paper7.pdf.

[28] D. Nagel, T. Affeldt, W. Balke, Data Narrations - Using flexible Data Bindings to support the Reproducibility of Claims in Digital Library Objects, in: Proceedings of the Workshop on Digital Infrastructures for Scholarly Content Objects (DISCO 2021) co-located with ACM/IEEE Joint Conference on Digital Libraries 2021(JCDL 2021), Online (Due to the Global Pandemic), September 30, 2021, volume 2976 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2021, pp. 19–23. URL: http://ceur-ws.org/Vol-2976/short-2.pdf.

[29] H. Kroll, J. Pirklbauer, J. Kalo, M. Kunz, J. Ruthmann, W. Balke, Narrative Query Graphs for Entity-Interaction-Aware Document Retrieval, in: Towards Open and Trustworthy Digital Societies - 23rd International Conference on Asia-Pacific Digital Libraries, ICADL 2021, Virtual Event, December 1-3, 2021, Proceedings, volume 13133 of *Lecture Notes in Computer Science*, Springer, 2021, pp. 80–95. doi:10.1007/978-3-030-91669-5_7.

[30] J. McCarthy, Notes on Formalizing Context, in: Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993, Morgan Kaufmann, 1993, pp. 555–562. URL: http://www-formal.stanford.edu/jmc/context3/context3.html.

[31] H. Kroll, J. Kalo, D. Nagel, S. Mennicke, W. Balke, Context-Compatible Information Fusion for Scientific Knowledge Graphs, in: Digital Libraries for Open Knowledge - 24th International Conference on Theory and Practice of Digital Libraries, TPDL 2020, Lyon, France, August 25-27, 2020, Proceedings, volume 12246 of *Lecture Notes in Computer Science*, Springer, 2020, pp. 33–47. doi:10.1007/978-3-030-54956-5_3.

[32] H. Kroll, F. Plötzky, J. Pirklbauer, W.-T. Balke, What a Publication Tells You—Benefits of Narrative Information Access in Digital Libraries, in: Proceedings of the 22nd ACM/IEEE Joint Conference on Digital Libraries, JCDL '22, ACM, New York, NY, USA, 2022. doi:10.1145/3529372.3530928.

[33] H. Kroll, N. Mainzer, W. Balke, On Dimensions of Plausibility for Narrative Information Access to Digital Libraries, in: Linking Theory and Practice of Digital Libraries - 26th International Conference on Theory and Practice of Digital Libraries, TPDL 2022, Padua, Italy, September 20-23, 2022, Proceedings, volume 13541 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 433–441. doi:10.1007/978-3-031-16802-4_43.

[34] F. Plötzky, W. Balke, What it Boils Down to... The Case for Event Analogies to Combat Com-

plexity in Digital Libraries, in: Proceedings of the Workshop on Digital Infrastructures for Scholarly Content Objects (DISCO 2021) co-located with ACM/IEEE Joint Conference on Digital Libraries 2021(JCDL 2021), Online (Due to the Global Pandemic), September 30, 2021, volume 2976 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2021, pp. 24–28. URL: http://ceur-ws.org/Vol-2976/short-3.pdf.

[35] F. Plötzky, W. Balke, It's the Same Old Story! Enriching Event-Centric Knowledge Graphs by Narrative Aspects, in: WebSci '22: 14th ACM Web Science Conference 2022, Barcelona, Spain, June 26 - 29, 2022, ACM, 2022, pp. 34–43. doi:10.1145/3501247.3531565.