

What a Publication Tells You – Benefits of Narrative Information Access in Digital Libraries

Hermann Kroll

kroll@ifis.cs.tu-bs.de

Institute for Information Systems, TU Braunschweig
Braunschweig, Lower Saxony, Germany

Jan Pirklbauer

j.pirklbauer@tu-bs.de

Institute for Information Systems, TU Braunschweig
Braunschweig, Lower Saxony, Germany

Florian Plötzky

ploetzky@ifis.cs.tu-bs.de

Institute for Information Systems, TU Braunschweig
Braunschweig, Lower Saxony, Germany

Wolf-Tilo Balke

balke@ifis.cs.tu-bs.de

Institute for Information Systems, TU Braunschweig
Braunschweig, Lower Saxony, Germany

ABSTRACT

Knowledge bases allow effective access paths in digital libraries. Here users can specify their information need as graph patterns for precise searches and structured overviews (by allowing variables in queries). But especially when considering textual sources that contain narrative information, i.e., short stories of interest, harvesting statements from them to construct knowledge bases may be a serious threat to the statements' validity. A piece of information, originally stated in a coherent line of arguments, could be used in a knowledge base query processing without considering its vital context conditions. And this can lead to invalid results. That is why we argue to move towards narrative information access by considering contexts in the query processing step. In this way digital libraries can allow users to query for narrative information and supply them with valid answers. In this paper we define narrative information access, demonstrate its benefits for Covid 19 related questions, and argue on the generalizability for other domains such as political sciences.

CCS CONCEPTS

• **Information systems** → **Information retrieval**; *Information integration*; **Web searching and information discovery**.

KEYWORDS

Narrative Information Access, Information Retrieval, Digital Libraries

ACM Reference Format:

Hermann Kroll, Florian Plötzky, Jan Pirklbauer, and Wolf-Tilo Balke. 2022. What a Publication Tells You – Benefits of Narrative Information Access in Digital Libraries. In *The ACM/IEEE Joint Conference on Digital Libraries in 2022 (JCDL '22)*, June 20–24, 2022, Cologne, Germany. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3529372.3530928>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

JCDL '22, June 20–24, 2022, Cologne, Germany

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9345-4/22/06...\$15.00

<https://doi.org/10.1145/3529372.3530928>

1 INTRODUCTION

From the beginnings of human language, knowledge was shared and passed on following a narrative oral tradition, i.e., they exchange stories and have structured debates and conversations [16]. With the advent of written language, these oral presentations were made persistent by writing up stories, comments and discussions in articles and books. The central way to encode all this knowledge is to tell a story: a narrator relates what was observed and how more complex conclusions were derived from basic claims. We thus understand this process as *composing narratives*, i.e., action patterns bound to real-world entities or concepts to form rich lines of arguments [7].

Today digital libraries play a key role in making knowledge publicly available in large-scale repositories. The necessary curation builds on a long-standing library sciences tradition and results in a variety of novel digital technologies to manage and access knowledge repositories, including the FAIR principles [27]: On the one hand, extensive collections need to be effectively maintained and efficiently archived. Here additional metadata enrichment is already used in each source to prepare the data for later access (Findability & Accessibility). On the other hand, digital libraries face an increasing amount of data collected from distributed sources. This is done either by providing unifying interfaces to individual collections of linked open data or by using information integration techniques over extractions from different sources (Interoperability & Reuse).

The traditional solution is to provide a simple keyword-based access path to the underlying data. Then users have to retrieve this data and *determine* what is actually *told* by the data. What happens here is that users try to *understand* the data to reuse the information of interest for their purposes. We understand this *exploratory process of understanding* as *gradually composing narratives*, in the sense of extracting and generalizing patterns that are 'told' by the data. Take, for instance, the COVID-19 pandemic: Patient records might describe suffered conditions after they have been vaccinated by a SARS-CoV-2 vaccine. Biomedical experts can then read through these records and extract typical story patterns, e.g., patients may experience headaches and pain, or even worse, may suffer from dangerous cerebral sinus venous thrombosis. Although this manual workflow is common, the rapid speed of the COVID-19 pandemics has shown that, given the amount of data available, it is hard to stay

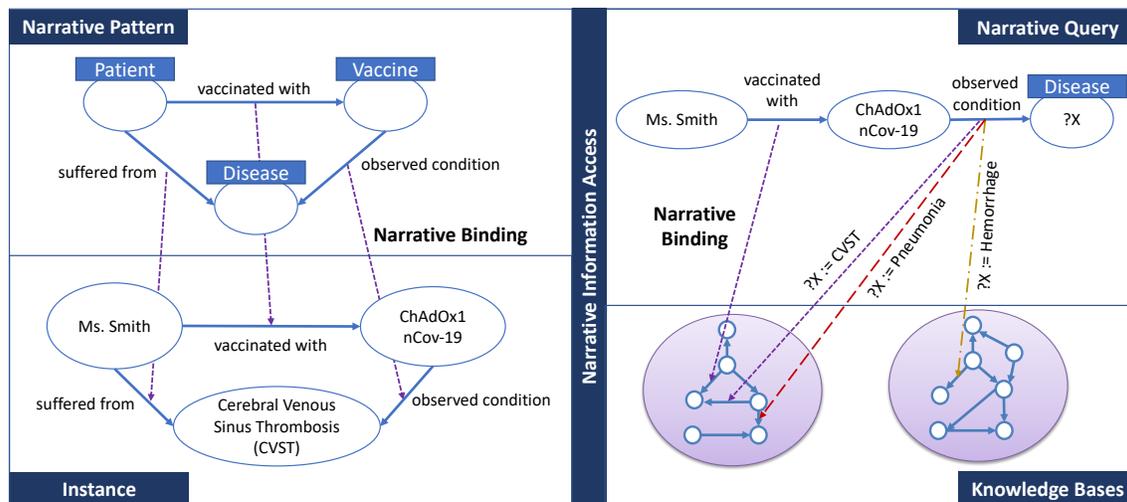


Figure 1: Systematic overview: A narrative pattern (upper left corner) describes a template how different entity types interact with each other. An instance then substitutes the entity types by concrete entities (lower left corner). These substitutions are called narrative bindings. On the right, the narrative query processing is depicted: Narrative bindings are found for each statement of a narrative query. Bindings that share the same context are depicted in the same colour and shape.

up-to-date. Even when restricting information sources only to well-curated ones, researchers would have to cover nearly 200k peer-reviewed articles about COVID-19 published in the US National Library of Medicine over the last two years¹.

Such rapid developments ask for novel and more efficient access methods. For example, a comprehensive database of all possible conditions observed in COVID-19 vaccinations might be helpful for improved diagnostics. Yet, when building such a knowledge base by harvesting statements about COVID-19 from textual sources, the answer quality may not be sufficient in practice. This is because the *observed conditions* are torn from the original course of vaccination as exhibited by some concrete patient. For example, some conditions might only be observed in elderly patients and thus, might not apply to children, or some complications might only be possible when a certain pre-existing condition is present in a patient. This means that although each condition was correctly extracted, the reusing of the resulting statements in a knowledge base may not be valid because the information's *contexts* do not match. When humans read through publications and retrieve arguments, they usually consider all essential context conditions such as the treated group or relevant pre-existing conditions. Moreover, in addition to contexts, humans also consider the connection between statements within a line of argument, e.g., do the assumptions within the arguments leading to a conclusion actually make sense together?

We argue that digital libraries need to move towards narrative information access, i.e., to offer query capabilities in the form of narrative patterns while considering vital contexts. Therefore we first define narrative information access. We then argue on contexts and how digital libraries can retain them. In addition, we perform two case studies on top of our narrative retrieval system, published last year [13]. We investigate COVID-19-related research questions in

cooperation with domain experts. We also asked an expert from the political sciences domain to study the system and describe how the political sciences domain could benefit from such a retrieval system. Finally, we discuss the generalizability, benefits, and challenges of narrative information access for digital libraries.

2 NARRATIVE INFORMATION ACCESS

In the following section we define the concept of narrative information access and discuss its key components. To ease understanding, we start with a running example from the biomedical field as a narrative pattern: Covid 19 vaccinations and their possible side effects. Consider the following short narrative:

EXAMPLE 1. *Some patients that were vaccinated by ChAdOx1 nCov-19 Vaccine (also known as Astra Zeneca) suffered Cerebral Venous Sinus Thrombosis (CVST). Hence Intracranial Sinus Thrombosis is an observed disease condition for the ChAdOx1 nCov-19 vaccine.*

Three types of entities participate in this example: a vaccine, patients, and a disease condition. In addition, three possible relations between the entity types are expressed: patients *are vaccinated* with the vaccine, patients *suffer from* a disease condition, and the disease condition *is observed* for the vaccine. Thus narrative patterns are described by typing their participants and naming their relations (see Fig. 1). The following ideas are based on an eased version of a narrative model that we introduced in [11].

Based on the encoding of knowledge in the Resource Description Framework (RDF) [18], we define narrative patterns by:

DEFINITION 1 (NARRATIVE PATTERN). *A narrative pattern is a connected, node- and edge-labeled directed graph, where each edge (labeled with a predicate name) represents a statement in the form of a (subject, predicate, object)-triple. Each node either represents a subject reflecting some entity type or an object reflecting either an entity type or literal values from a certain domain.*

¹<https://www.ncbi.nlm.nih.gov/research/coronavirus/>

Any knowledge base in RDF format can then be seen as a graph containing a collection of *instances of* narrative patterns as sub-graphs, i.e., all nodes have been instantiated (either by URIs in the case of entities or by concrete literal values). We can now translate our previous example narrative using a narrative pattern as a kind of skeleton for the narrative. A possible instance is depicted in Fig. 1 (please note that for simplification, we replaced long URI prefixes with short entity names).

In brief, we have a graph representation of a concrete narrative structured by some narrative pattern. Hence narrative patterns can be understood as (sub-)graphs isomorphisms on RDF knowledge bases. We then define narrative queries using such patterns:

DEFINITION 2 (NARRATIVE QUERY). *A narrative query is a narrative pattern where each node is either instantiated by a concrete entity or literal value or replaced by a variable (labeled by a variable name).*

By design our proposed querying method has very similar semantics to querying RDF knowledge bases with SPARQL: If a narrative query does not contain a variable, then the answer is whether there exists an instance in the knowledge base that is isomorphic to the query's narrative pattern and features all the query's exact entities/literal values in the right places (cf. ASK queries in SPARQL). If a narrative query contains one or more variables, then these variables must be substituted by concrete entities from the knowledge base during query processing. Of course, all matches to the query must be valid with regard to variable substitutions, i.e., the substituted pattern and the respective entities/values must be contained in the knowledge base. We understand such a matching process as *binding a query* [12], i.e., we take some edge of the query's narrative pattern and bind it against a knowledge base edge and bind concrete entities and literal values to the respective entity types or literal domains in the pattern.

Returning to our example, we may query which *disease conditions* the *ChAdOx1 nCov-19* vaccinated patient *Smith* could possibly suffer from. The respective narrative query is depicted in Fig. 1. The first step to answer this query is to compute narrative bindings against the underlying knowledge base(s). We may find a binding b_1 confirming that *Ms. Smith* has been vaccinated with *ChAdOx1 nCov-19*. In addition, we must *substitute* the variable *?X (of type disease)*. Here we may find three bindings with suitable substitutions: b_2 (CVST), b_3 (Pneumonia), and b_4 (Hemorrhage). In common graph querying we would now join the intermediate results to list all conditions that *Ms. Smith* could possibly expect: CVST, pneumonia and hemorrhage.

Now, assume for the time being that *pneumonia* have only been observed in elderly people, whereas *Ms. Smith* is still young. Then *pneumonia* as a possible side effect of the vaccination might no longer apply to *Ms. Smith*, although the respective binding observing *pneumonia* as a possible side effect of a *ChAdOx1 nCov-19* vaccination is perfectly correct. The problem here is that b_3 would not be valid *in general*, because the observed conditions do not apply to all patients, but only to elderly patients. Although the bindings are correctly retrieved, not all of them might actually fit into the context of *Ms. Smith*.

Here information was torn apart regarding a sensitive context such as the target group information. One might argue that extracting RDF-style knowledge from individual patient records could even in the best case be problematic and should not be done in this way. While we agree that all patients are somewhat unique cases, this kind of extraction is common practice in real life applications, e.g., the *causes* relation in SemMedDB [8], *medical causes* in Wikidata [26]², and *causes* in DBpedia [1]³.

The effect is that even if knowledge bases did only contain correct statements, fusing them to answer a query may still produce incorrect results. Indeed, it is a good scientific practice to arrange statements as complex lines of arguments, i.e., authors are sure to mention all essential contexts, settings, assumptions made, necessary conditions, hypotheses, experimental designs, etc. It is essential to fuse only those arguments fitting into the same context provided in the form of constraints by other arguments or the query terms. We call bindings *context-compatible* if they can safely be fused to form valid knowledge. Based on the idea of context-compatibility, we are now ready to propose a novel query processing method that considers contexts as constraints upon the query process to bypass the previous issues.

DEFINITION 3 (NARRATIVE QUERY PROCESSING). *Given a narrative query and a set of knowledge bases, the query processing has to a) bind each individual query statement against underlying data of the knowledge base(s) and b) check the context-compatibility of the computed bindings. The result of the query process is thus a set of valid bindings, individually binding all query statements and being context-compatible.*

Thus narrative query processing ensures that contexts are considered while matching graph patterns. All bindings must in this way share a compatible context. And with this narrative query processing method we can now define narrative information access:

DEFINITION 4 (NARRATIVE INFORMATION ACCESS). *Narrative Information Access allows users to formulate their information need as a narrative query. A narrative retrieval system then performs narrative query processing for this pattern and returns the results to the user. If results are found, we call the narrative pattern plausible.*

2.1 The Problem of Context-Compatibility

In this section we investigate the problem of context-compatibility in more detail and discuss suitable solutions how digital libraries can retain contexts in practice. Contexts define the *scope* in which a piece of information can be fused with other statements. This means that a context has to involve all information that need to be known to validate some larger, fused piece of information. But unfortunately, essential parts of contexts may get lost during information extraction. Generally speaking, problems with context compatibility come in at least two distinct flavors: *constraining contexts and correspondence contexts*. Constraining contexts scope the validity of fusions of statements over the entire query, i.e., for some statements in a substitution, a fusion is impossible because they have been extracted from contradicting contexts. In contrast, correspondence contexts limit the actual fusion of individual pieces

²<https://www.wikidata.org/wiki/Property:P828>

³<https://dbpedia.org/property/causes>

of knowledge between which a fusion would generally be possible but is not warranted by the data from which the information was extracted.

For a problematic case with *constraining contexts* consider the following example:

EXAMPLE 2. “We report a case of a 62-year-old man who developed cerebral venous sinus thrombosis with subarachnoid hemorrhage and concomitant thrombocytopenia, which occurred 13 days after ChAdOx1 nCov-19 injection.” [2]

Among others we may extract the following statements:

- (patient, vaccinated by, ChAdOx1 nCov-19)
- (patient, suffered from, cerebral venous sinus thrombosis)

But the statement that some patient suffered from cerebral venous sinus thrombosis is only sensible within the context of this particular patient record. Unfortunately, there is no information whether the statement can be generalized to other patients. Thus if the extractions’ context (e.g., the patient’s age, or that he was recently vaccinated) is lost, information fusions or reasoning processes relying on this specific piece of information may produce invalid results and even run into inconsistencies.

In brief, constructing knowledge bases with insufficiently contextualized statements and then using them to answer complex query patterns may result in invalid answers: Vaccinations with ChAdOx1 nCov-19 may indeed lead to a *pneumonia* although probably not in all contexts.

For a problematic case with *corresponding contexts* consider the following example:

EXAMPLE 3. “Secondary analyses found increased risk of CVST after ChAdOx1 nCov-19 vaccination (4.01, 2.08 to 7.71 at 8-14 days), after BNT162b2 mRNA vaccination (3.58, 1.39 to 9.27 at 15-21 days), and after a positive SARS-CoV-2 test.” [9]

We may extract the following statements:

- (ChAdOx1 nCov-19, observed condition, CVST)
- (BNT162 Vaccine, observed condition, CVST)
- (CVST, risk after vaccination, 4.01)
- (CVST, risk after vaccination, 3.58)

Now information fusion for answering the query (*?x, observed condition, CVST*) AND (*CVST, risk after vaccination, ?y*). would compute the Cartesian product producing four results (two of which are correct, while the other two are incorrect). This is because the binary extraction has lost the information, which risk factor belongs to which vaccine.

In brief, although all statements are mentioned within the close scope of a clinical trial having inclusion and exclusion criteria, an information extraction process may lose how statements belong together within that context.

Here the text expresses a ternary relation between *vaccines*, *conditions* and *probabilities* that is broken down into binary relations. Moreover, note that this is not an artifact of automatic processes, as even manual extraction may yield the same result because of the restriction of using only binary relations.

In conclusion, although all of our example statements were syntactically correct, vital semantics have been lost because the context was neglected. This forms a serious threat to the **validity** of query

results, i.e., even correctly extracted but subsequently fused statements may not always produce valid answers in query processing or reasoning. Specifically, invalid answers are those cases that do not match the user’s context or connect statements that do not belong together.

Since these problems are the main reason we argue to move towards narrative information access, we will take a closer look at possible remedies in the following section.

2.2 Maintaining Contexts in Digital Libraries

So how can we retain contexts in practical digital library projects? This subsection discusses research and methods to combat both loss of constraining contexts and loss of correspondence contexts.

N-ary Relations. Ernst et al. [6] proposed an n-ary extraction method to precisely retain complex relations, e.g., a relation *vaccinated_patients_suffer* that involves the *target group*, *vaccine* and *side effects*. However designing appropriate n-ary relation signatures a-priori is challenging because it requires extensive domain knowledge. The authors collected examples to train a suitable extraction model for their relations. In addition, they performed partial reasoning to compose partial statements to n-ary statements because their extraction method was also limited to sentences. The reasoning step helped to increase the extraction recall but required the definition of rules (which facts should be composed). Although n-ary relations are strongly appreciated, practical extraction methods hardly support them because defining signatures, providing enough training examples, and formulating reasoning constraints is an exhausting task.

Explicit Context Models. McCarthy introduced an explicit context model based on the first-order predicate logic [19]. The model allows users to formulate context conditions for arbitrary statements explicitly. In addition, he discussed relations between contexts, e.g., one context might *specialize* another context. Hand-crafted rules were then formulated to determine how to combine contexts and their enclosed statements. *VIKEF* is an example digital library project supporting explicit context information in an RDF knowledge base [23].

Implicit Contexts. We proposed using document references as an implicit and practical context model [10]. We suggested to store references to the source documents when harvesting statements from it. These references were then used to estimate which statements can safely be combined to produce valid answers. When combining only statements extracted from the same document, the resulting precision in a downstream application will increase, but the recall is bound to decrease. We therefore proposed measures to estimate *compatibility* between contexts to flexibly manage the precision/recall trade-off, e.g., text and author similarities.

Such implicit context models might be suitable candidates to retain context in digital libraries because they are cheap to maintain, i.e., only references to the statements’ sources must be retained. But their quality and explainability are somewhat limited, e.g., how should we explain why two documents are context-compatible based on some text similarity measure. Keyword extraction might be a good method to retrieve context proxies here; See YAKE [5] for example. In summary, implicit context models are easy to use

and may yield good precision, but estimating context-compatibility remains challenging, and the overall quality achieved might still not be good enough for digital libraries.

Provenance. Provenance information is often understood to be any kind of information that may validate some statement's quality or origin [28]. Provenance might range from storing a reference to the statement's origin to storing information about the creation process, e.g., author, release date, point in time, and more. The Prov-O Ontology Description is a common standard for defining and storing general provenance information [17]. Prov-O supports complex provenance graphs to describe the origin of some statements. As an alternative, the Wikidata project supports qualifiers (property-value pairs) to retain provenance for its statements [26], e.g., references, determination methods, time and location information.

Nevertheless, using qualifiers and provenance information in practical applications, especially in query processing, remains an exception. Returning to our example, how could we use a qualifier information about the *62-year-old man* in query processing? Should we formulate hand-crafted rules on how different provenance information affects the actual query processing? How do we know when qualifiers describe the same or a compatible context? We understand Prov-O and provenance in general as possible implementations to store contexts. However they do not provide a ready-to-use solution to retain both by default. Domain experts and digital library curators must carefully define corresponding statements and describe how they are used for a practical application.

3 NARRATIVE QUERY PROCESSING IN PRACTICE – CASE STUDIES

We performed case studies to understand the benefits and limitations of narrative information access. In particular, we built on our publicly available narrative retrieval system called *Narrative Query Graphs for Entity-Interaction Document Retrieval* by [13]. We built a working document retrieval system that allows formulating information needs as graph patterns, i.e., entities and their corresponding interactions. We transformed biomedical document abstracts into a graph representation called document graph as knowledge bases. Then the retrieval system allows matching user queries against these document graphs and returns all matches. Since document graphs match queries only within single documents, contexts are to some degree considered in query processing because the context can quite safely be assumed to be consistent within each document abstract.

3.1 Narrative Query Graphs for Covid 19

In cooperation with pharmaceutical domain experts, the Robert-Koch Institute in Germany and the ZB MED library, we enhanced the narrative retrieval system to answer Covid 19-related research questions:

- (1) We included the LitCovid collection from PubMed (peer-reviewed articles about Covid 19) and the latest Covid 19-related pre-prints supplied by ZB MED [14, 15]. These pre-prints can be accessed via their Preview service⁴.

- (2) We developed a vaccine entity vocabulary by utilizing Wikidata and the Medical Subject Headings (MeSH). In addition, we derived an entity for Long Covid 19 from MeSH.

The prototype of the enhanced narrative query system is publicly available⁵. In the following we investigate whether typical research questions from the pharmacy domain can be translated into narrative query graphs and how helpful such searches are in practice. Please note that this case study does not yet contain a comprehensive evaluation. We are currently preparing a large-scale study with our partners.

Long Covid Related Questions. The development of the Covid 19 pandemics has shown that *Long Covid* is a severe threat to a patient's health. So what are *common symptoms that are reported for Long Covid*? We formulated the following query graph: (*post-acute COVID-19 syndrome, associated, ?X(Disease)*). *?X(Disease)* means that we search with a variable named *?X* that should be substituted by entities of the type *Disease*. *Post-acute COVID-19 syndrome* is an entity from the Medical Subject Headings (MeSH)⁶. The system responded with a list of commonly known conditions such as *Fatigue* (44), *Dyspnea* (19), *Anosmia* (10), *Cognitive Dysfunction* (9) and *Headache* (7). The number in brackets refers to how many documents share the corresponding variable substitution. The system can show the origin of the extraction, i.e., the sentence in which the pattern was matched. However also substitutions such as *Covid 19* (143) and *Infections* (61) were not helpful.

We adjusted the previous query to search for patient cases: (*post-acute COVID-19 syndrome, associated, Human*) AND (*Human, associated, ?X(Disease)*). Here *Humans* is an entity that stand for *patients, men, women*, etc. The current version of the system did not support searching for specific target groups. This query could be matched against abstracts such as: “[...] *post-COVID-19 syndrome in patients with primary Sjogren's syndrome (pSS) affected by acute SARS-CoV-2 infection. [...] More than 40% of pSS patients reported the persistence of four symptoms or more, including anxiety/depression (59%), arthralgias (56%), sleep disorder (44%), fatigue (40%), anosmia (34%) and myalgias (32%).*” [3] Here the implicit context ensured that both statements must be matched against a single abstract. But the number of found results were decreased: *Fatigue* (15), *Dyspnea* (8), *Cognitive Dysfunction* (4) and *Headache* (3).

A quick look over both results revealed that publications were missed because they did not explicitly contain the entity *post-acute COVID-19 syndrome*. Instead, publications may describe Covid 19 infections and observations made six months later. Here entity linking did not detect the explicit entity.

Vaccinations. We formulated a query to list commonly used vaccines that are associated with Covid 19: (*Covid 19, associated, ?X(Vaccine)*). Helpful substitutions were for example: *BTN162 aka Pfizer* (175), *ChAdOx1 nCoV-19 aka Astra Zeneca* (79), and *2019-nCoV Vaccine mRNA-1273 aka Moderna* (76). In addition, miss leading substitutions like *Vaccine* (3472) and *Covid-19 Vaccines* (685) were found and not helpful because they were far too general. We enhanced the query by asking for common side effects of *ChAdOx1 nCoV-19*: (*ChAdOx1 nCoV-19, associated, ?X(Disease)*). Substitutions such as

⁴<https://preview.zbmed.de/>

⁵<http://www.pubpharm.de/services/prototypes/narratives/>

⁶<https://meshb.nlm.nih.gov/record/ui?ui=C000711409>

Thrombosis (93), *Thrombocytopenia* (79), and *CVST* (18) were found. The system yielded also not helpful results like *Covid-19* (79) and *Infections* (27) caused by wrong extractions. Again, we added the *Human* entity to precisely query for studies: (*Human*, associated, ?X(*Disease*)) AND (*Chadox1 Ncov-19*, associated, *Human*). Here we could quickly find a case study [25] for CVST investigation.

Treatments. We were also interested in queries that consider treatments for Covid-19 symptoms. Therefore, we formulated the query: (?X(*Drug*), *treats*, *Covid 19*). Helpful substitutions were *Hydroxychloroquine* (829) and *Remdesivir* (581). The system’s provenance information (matched sentences) showed that the system found the statement in sentences like: “An example of which is *remdesivir* which has now been approved for use in COVID-19 patients by the US Food and Drug Administration.” [4] We rewrote the query by integrating the patient again, similar to the previous approaches. Here we retrieved matches such as “We identified 55 patients who were treated with *remdesivir* for COVID-19 and analyzed inflammatory markers and clinical outcomes.” [22]

Discussion. The case study showed that narrative information access indeed could support typical tasks like generating structured overviews of the latest literature or quickly finding precise hits: On the one hand, suitable substitutions for *Long Covid 19 symptoms* or *Covid 19 drug treatments* were indeed found, thus successfully structuring the latest literature. On the other hand, the expressive query format enabled the integration of *patients* in the query to ensure that the results had to connect the disease or drug to a concrete target group.

As a small caveat, note that all queries were matched only against implicit document contexts, ensuring the statements’ context compatibility. In this way retaining the context for query processing came cheap: The origin of the statements needed to be stored and the query processing had to be restricted to document graphs. Of course, this (overly careful) restriction to document graphs also comes with severe limitations since combining knowledge from different sources is a common practice and vital necessity in scientific research. While the precision in our query tasks was very high and thus matches were accurate, the respective recall was admittedly marginal. More open yet effective measures for controlling context-compatibility than using documents graphs will be needed to build large-scale practical narrative retrieval systems (as previously discussed in section 2.2).

3.2 Narrative Query Graphs in Political Sciences

In cooperation with the specialized information service for political sciences [21](Pollux)⁷ we were interested how the political sciences can benefit from narrative information access. We asked an expert (Ph.D. in political sciences) to study the biomedical narrative query graph retrieval system. He then formulated questions that would be of interest in political sciences. Due to the lack of available knowledge bases we could not realize a practical retrieval system here. Instead, we went through two of his questions and argue in the following how they could be answered and why narrative information access is vital. In addition we report on opportunities

and potential obstacles in political sciences. In the following we picked two of his questions as showcases:

- (1) *How do heads of government in Latin America and Scandinavia present the question what action is needed in relation to climate change?*
- (2) *How do Germany’s major daily newspapers negotiate the course of the refugee crisis in 2015 and 2016?*

So why do we need narrative information access to answer his questions? The main reason here is that both questions asked to combine several information: For the first question, we have to combine statements about climate change in the time period of corresponding presidents (temporal and location context). The temporal and location contexts and the source of information (the heads of government) are vital to determine statements’ validity. For the second question, we have to generate a structured overview of statements and viewpoints (e.g., conservative, progressive, etc.) from daily newspapers about the refugee crisis in 2015 and 2016 (temporal context, framing, and wording). The selection of keywords (wording) may express different viewpoints. Again, context (e.g., the kind and target group of a newspaper) was vital to align the statements with a certain viewpoint.

Parts of both queries could be answered with today’s knowledge bases already. Consider, for example, the usage of Wikidata: Concerning question (1), formulating a SPARQL query allowed us to retrieve a list of heads of governments in both geographical regions. And we could also combine the results with their temporal context:

- (?country, head_of_state, ?stmt) AND
(?stmt, head_of_state, ?person) AND
(?stmt, start_time, ?time) AND
(?country, part_of, Latin America).

Note, the *?stmt* notation is necessary to query Wikidata for qualifiers. This query yielded 66 results.

Concerning question (2), major newspaper could be easily identified by querying Wikidata: (?newspaper, instance_of, daily newspaper) AND (?newspaper, country, Germany). Querying Wikidata resulted in 58 newspapers. Newspapers are often associated with a *political ideology*. And indeed, Wikidata stores information that the *Frankfurter Allgemeine Zeitung (FAZ)* has the *political ideology liberal conservatism*⁸. In this way we could derive additional context information when analyzing statements from a newspaper. Note that this might be a good approximation but newspapers might also include articles that follow different ideologies.

The next part would include context-sensitive information retrieval based on the Wikidata results. To answer both questions, we had to rely on texts, e.g., from Pollux or specialized knowledge bases for claims such as ClaimsKG [24]. Here a comprehensive extraction is necessary to identify statements in texts.

But even if a knowledge base had been available, question (2) asked for different levels of granularity regarding the context of statements. In a simple scenario, it might be enough to extract statements from news articles and cluster them by their *political ideology* from Wikidata if available. However guest commentary or changes in the editorial board might include statements that stemmed from a different ideology. Therefore, we have to classify

⁷<https://www.pollux-fid.de>

⁸<https://www.wikidata.org/wiki/Q10184>

the ideology based on an article's wording and framing, and may not solely rely on the general ideology of a newspaper.

Challenges. Political sciences have a broad range of essential concepts, e.g., viewpoints, schools of thought, and ambiguous terms. These concepts are hard to identify in a text, unlike biomedical entities. Here wording and framing of texts might determine the viewpoint, whereas a drug in medicine remains the same drug regardless of wording. Moreover, central terms like "Democracy" or "Society" are not unambiguously defined and can be interpreted differently, depending on a school of thought. Furthermore, even if we identify the concepts, extracting structured information remains challenging. Statements in this domain are more complex than just expressing a binary relation between a patient and a disease condition.

These issues have to be addressed to realize a convenient narrative information access. Although solving them remains challenging, the previous cases showed that political sciences could benefit from such access. Structuring publications into schools of thought or clustering viewpoints regarding a topic would be beneficial here. Moreover, without considering the context of information, such access could hardly be realized.

3.3 Investigating Common Knowledge Bases

After we performed both case studies, we also were interested in the generalizability of the benefits of narrative information access to other domains. We first had a look at publicly available knowledge bases for their application and possible issues.

Interestingly, the following statement is included in Wikidata⁹:

- (*Barack Obama, born in, Kenya*)

In Wikidata this statement is complemented by a qualifier that states *mentioned in a conspiracy theory*. A qualifier is a statement about some other statement, i.e., a property-value pair attached to a statement. But this incorrect statement that *Barack Obama was born in Kenya* can only be sensible when considering it in the context of some *conspiracy theory*. Wikidata marks this data in their user interface by a colour encoding: *green* for fact-checked and *red* for not fact-checked. However the decision whether something is fact-checked or not is often not easy, e.g., partially fact-checked statements. In addition, different school of thoughts may accept or reject a certain statement. And having a general decision here, whether something is *true* or not, remains open.

We found another interesting example in the real-world knowledge base DBpedia¹⁰.

- (*Barack Obama, was, Senator of Illinois*)
- (*Barack Obama, predecessor, Peter G. Fitzgerald*)
- (*Barack Obama, was, U.S. President*)
- (*Barack Obama, predecessor, George W. Bush*)

Suppose a user asks the following query: *Who was the predecessor of the U.S. President Barack Obama?* In that case the results are *George W. Bush* (correct) and *Peter G. Fitzgerald* (wrong). Thus querying DBpedia with such queries can lead to wrong results. The example query could have been answered correctly if the connection between the statements had been retained.

⁹<https://www.wikidata.org/wiki/Q76> 01.2022

¹⁰https://dbpedia.org/page/Barack_Obama 01.2022

Both examples show that the loss of context is also an issue in common knowledge bases. Information can quickly be broken down lossy and cannot be reassembled lossless afterward.

4 DISCUSSION

Narrative information access ensures that the binding process must consider contexts when making a narrative plausible. Here bindings must be context-compatible which ensures that the bindings form valid answers. We do not claim that knowledge bases cannot do the job. But if they are built without considering context and statements are restricted to triples, then information is broken down in a lossy fashion and cannot be reassembled lossless afterward. Thus contexts definitely have to be considered when designing knowledge bases to supply narrative information access.

4.1 Generalizability

Although we made our central use case in the biomedical domain, we argue that we can generalize our findings across domains. The Obama examples show how easily context can be lost in common knowledge bases. In addition, we reported on opportunities and challenges in political sciences. Here proposed use cases showed how beneficial narrative information access could be. Due to the lack of structured knowledge bases, we could hardly realize an access here. But context like temporal periods or a newspaper's viewpoint is essential to answer narrative queries correctly.

4.2 Benefits for Digital Libraries

The Covid 19 pandemics has shown how important it is to carefully handle scientific claims. Tearing such claims apart from the original lines of arguments has caused many miss leading debates (based on fake news) and movements across the world. Digital libraries should head for a more comprehensive knowledge curation by allowing narrative information access. Here the vital contexts are considered when answering queries. Our case study has shown how context-aware query systems can be applied to Covid 19 related questions. Although our study lacked a comprehensive evaluation, we demonstrate such benefits in practice: Narrative Information access allows to structure the latest literature or quickly find suitable information. Realizing and implementing suitable workflows may be cost-intensive, but digital libraries can benefit from them.

4.3 Future Work

A new challenge that has to be addressed for narrative information access is the growing heterogeneity of data sources with digital libraries, such as textual sources, image collections, experimental data or structured knowledge bases. Research data sets are a good consideration to link narrative queries against [20]. Making these heterogeneous repositories accessible in a unified way and integrating their different kinds of information requires effective access paths that often have to be intelligently customized to the content types. For narrative information access this means that bindings on (sub-)graphs of narrative queries have to be computed against extractions (either precomputed or extracted on-the-fly) from different media. Investigating such extraction is thus essential for broader applicability of narrative retrieval systems.

5 CONCLUSION

Although knowledge bases allow effective access paths in digital libraries, we demonstrated their limitations when handling narrative information. Here information, originally stated in coherent lines of arguments, can be broken into pieces that cannot be reassembled lossless afterward. This paper defines narrative information access as an extension to common knowledge base querying. Here the context of statements must be retained and considered to produce valid answers when querying narrative information. Realizing narrative information access in digital libraries can be cost-intensive in practice, but like the case study for Covid 19 retrieval has shown, implicit document contexts may approximate it. The examples of Barack Obama in common knowledge bases, our investigation in Covid 19 related questions, and the discussion in political sciences have shown how beneficial narrative information access can be. Even now existing methods and techniques can be used to implement narrative information access in digital libraries reliably. However handling heterogeneous library content (research data, tables, images, etc.) would be the next step to enhance such access further.

ACKNOWLEDGMENT

Supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation): PubPharm – the Specialized Information Service for Pharmacy (Gepris 267140244).

REFERENCES

- [1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *The semantic web*. Springer, Busan, Korea, 722–735.
- [2] Alice Bérezné, David Bougon, Florence Blanc-Jouvan, Nicolas Gendron, Cecile Janssen, Michel Muller, Sébastien Bertil, Florence Desvard, Isabelle Presot, Benjamin Terrier, et al. 2021. Deterioration of vaccine-induced immune thrombotic thrombocytopenia treated by heparin and platelet transfusion: Insight from functional cytometry and serotonin release assay. *Research and Practice in Thrombosis and Haemostasis* 5, 6 (2021), e12572.
- [3] P. Brito-Zerón, N. Acar-Denizli, V. C. Romão, B. Armagan, R. Seror, F. Carubbi, S. Melchor, R. Priori, V. Valim, S. Retamozo, S. G. Pasoto, V. F. M. Trevisani, B. Hofauer, A. Szántó, N. Inanc, G. Hernández-Molina, A. Sebastian, E. Bartoloni, V. Devauchelle-Pensec, M. Akasbi, F. Giardina, M. Bandeira, A. Sisó-Almirall, and M. Ramos-Casals. 2021. Post-COVID-19 syndrome in patients with primary Sjögren’s syndrome after acute SARS-CoV-2 infection. *Clin Exp Rheumatol* 39 Suppl 133, 6 (2021), 57–65.
- [4] A. B. Butnariu, A. Look, M. Grillo, T. A. Tabish, M. J. McGarvey, and M. Z. I. Pranjoli. 2022. SARS-CoV-2-host cell surface interactions and potential antiviral therapies. *Interface Focus* 12, 1 (Feb 2022), 20200081.
- [5] Ricardo Campos, Vitor Mangaravite, Arian Pasquali, Alípio Jorge, Célia Nunes, and Adam Jatowt. 2020. YAKE! Keyword extraction from single documents using multiple local features. *Information Sciences* 509 (2020), 257–289. <https://doi.org/10.1016/j.ins.2019.09.013>
- [6] Patrick Ernst, Amy Siu, and Gerhard Weikum. 2018. HighLife: Higher-Arity Fact Harvesting. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) (*WWW ’18*). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1013–1022. <https://doi.org/10.1145/3178876.3186000>
- [7] James B Freeman. 2011. *Argument Structure: Representation and Theory*. Vol. 18. Springer Science & Business Media, Berlin/Heidelberg, Germany.
- [8] Halil Kilicoglu, Dongwook Shin, Marcelo Fiszman, Graciela Rosemblat, and Thomas C. Rindfleisch. 2012. SemMedDB: a PubMed-scale repository of biomedical semantic predications. *Bioinformatics* 28, 23 (10 2012), 3158–3160. <https://doi.org/10.1093/bioinformatics/bts591>
- [9] O. H. Klungel and A. Pottegård. 2021. Strengthening international surveillance of vaccine safety. *BMJ* 374 (08 2021), n1994.
- [10] Hermann Kroll, Jan-Christoph Kalo, Denis Nagel, Stephan Mennicke, and Wolf-Tilo Balke. 2020. Context-Compatible Information Fusion for Scientific Knowledge Graphs. In *Digital Libraries for Open Knowledge*. Springer, Lyon, France, 33–47. https://doi.org/10.1007/978-3-030-54956-5_3
- [11] Hermann Kroll, Denis Nagel, and Wolf-Tilo Balke. 2020. Modeling Narrative Structures in Logical Overlays on Top of Knowledge Repositories. In *Conceptual Modeling*. Springer, Vienna, Austria, 250–260. https://doi.org/10.1007/978-3-030-62522-1_18
- [12] Hermann Kroll, Denis Nagel, Morris Kunz, and Wolf-Tilo Balke. 2021. Demonstrating Narrative Bindings: Linking Discourses to Knowledge Repositories. In *Fourth Workshop on Narrative Extraction From Texts, Text2Story@ECIR2021 (CEUR Workshop Proceedings, Vol. 2860)*. CEUR-WS.org, 57–63. <http://ceur-ws.org/Vol-2860/paper7.pdf>
- [13] Hermann Kroll, Jan Pirklbauer, Jan-Christoph Kalo, Morris Kunz, Johannes Ruthmann, and Wolf-Tilo Balke. 2021. 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 (Lecture Notes in Computer Science, Vol. 13133)*. Springer, Online, 80–95. https://doi.org/10.1007/978-3-030-91669-5_7
- [14] Lisa Langnickel, Roman Baum, Johannes Darms, Sumit Madan, and Juliane Fluck. 2021. COVID-19 preVIEW: Semantic Search to Explore COVID-19 Research Preprints. In *Public Health and Informatics*. IOS Press, Amsterdam, the Netherlands, 78–82. <https://doi.org/10.3233/SHTI210124>
- [15] Lisa Langnickel, Johannes Darms, Roman Baum, and Juliane Fluck. 2021. preVIEW: from a fast prototype towards a sustainable semantic search system for central access to COVID-19 preprints. *Journal of EAHIL* 17, 3 (Sep. 2021), 8–14. <https://doi.org/10.32384/jeahil17484>
- [16] János László. 2008. *The science of stories: An introduction to narrative psychology*. Routledge, Oxfordshire, England, UK.
- [17] T. Lebo, S. Sahoo, and D. McGuinness. 2013. PROV-O: The PROV Ontology. <https://www.w3.org/TR/prov-o/>.
- [18] Frank Manola, Eric Miller, Brian McBride, et al. 2004. RDF primer. *W3C recommendation* 10, 1-107 (2004), 6.
- [19] John McCarthy. 1993. Notes on Formalizing Context. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence, Chambéry, France, August 28 - September 3, 1993*. Morgan Kaufmann, Chambéry, France, 555–562. <http://www-formal.stanford.edu/jmc/context3/context3.html>
- [20] Denis Nagel, Till Affeldt, and Wolf-Tilo Balke. 2021. 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 (CEUR Workshop Proceedings, Vol. 2976)*. CEUR-WS.org, Online, 19–23. <http://ceur-ws.org/Vol-2976/short-2.pdf>
- [21] Tim Schardelmann and Wolfgang Otto. 2018. POLLUX – von der Bedarfsanalyse zur technischen Umsetzung. *Bibliotheksdienst* 52, 3-4 (2018), 225–234. <https://doi.org/10.1515/bd-2018-0029>
- [22] K. Stoeckle, B. Witting, S. Kapadia, A. An, and K. Marks. 2022. Elevated inflammatory markers are associated with poor outcomes in COVID-19 patients treated with remdesivir. *J Med Virol* 94, 1 (01 2022), 384–387.
- [23] Heiko Stoermer, Ignazio Palmisano, Domenico Redavid, Luigi Iannone, Paolo Bouquet, and Giovanni Semeraro. 2006. Contextualization of a RDF Knowledge Base in the VIKEF Project. In *Digital Libraries: Achievements, Challenges and Opportunities*. Springer Berlin Heidelberg, Berlin, Heidelberg, 101–110.
- [24] Andon Tchechmedjiev, Pavlos Fafalios, Katarina Boland, Malo Gasquet, Matthäus Zloch, Benjamin Zapilko, Stefan Dietze, and Konstantin Todorov. 2019. ClaimsKG: A Knowledge Graph of Fact-Checked Claims. In *The Semantic Web – ISWC 2019*. Springer International Publishing, Cham, 309–324.
- [25] C. Thompson, H. Karunadasa, D. Varma, M. Schoenwaelder, and W. Clements. 2021. Impact of COVID vaccination rollout on the use of computed tomography venography for the assessment of cerebral venous sinus thrombosis. *J Med Imaging Radiat Oncol* 65, 7 (Dec 2021), 883–887.
- [26] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (2014), 78–85.
- [27] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* 3, 1 (15 Mar 2016), 160018. <https://doi.org/10.1038/sdata.2016.18>
- [28] M. Wylot, P. Cudré-Mauroux, M. Hauswirth, and P. Groth. 2017. Storing, Tracking, and Querying Provenance in Linked Data. *IEEE Transactions on Knowledge and Data Engineering* 29, 8 (2017), 1751–1764.