

# Gravsearch: transforming SPARQL to query humanities data

Tobias Schweizer<sup>a,b</sup> and Benjamin Geer<sup>a,c</sup>

<sup>a</sup> *Data and Service Center for the Humanities, Gewerbestrasse 24, 4123 Allschwil, Switzerland*

<sup>b</sup> *E-mail: tobias.schweizer@dasch.swiss*

<sup>c</sup> *E-mail: benjamin.geer@dasch.swiss*

**Editors:** Antonis Bikakis, University College London, UK; Beatrice Markhoff, University of Tours, FR; Alessandro Mosca, Faculty of Computer Science, Free University of Bozen-Bolzano, IT; Stephane Jean, University of Poitiers - ENSMA, FR; Eero Hyvönen, University of Helsinki, Aalto University, Finland

**Solicited reviews:** Pietro Liuzzo, Hiob Ludolf Centre for Ethiopian Studies, Universität Hamburg, Germany; Martin Rezk, DMM.com, Japan; Benjamin Cogrel, KRDB, Free University of Bozen-Bolzano, Italy

**Abstract.** RDF triplestores have become an appealing option for storing and publishing humanities data, but available technologies for querying this data have drawbacks that make them unsuitable for many applications. Gravsearch (Virtual Graph Search), a SPARQL transformer developed as part of a web-based API, is designed to support complex searches that are desirable in humanities research, while avoiding these disadvantages. It does this by introducing server software that mediates between the client and the triplestore, transforming an input SPARQL query into one or more queries executed by the triplestore. This design suggests a practical way to go beyond some limitations of the ways that RDF data has generally been made available.

**Keywords:** SPARQL, humanities, querying, qualitative data, API

## 1. Introduction

Gravsearch transforms SPARQL queries and results to facilitate the use of humanities data stored as RDF.<sup>1</sup> Efforts have been made to use RDF triplestores for the storage and publication of humanities data [1–3], but there is a lack of appropriate technologies for searching this data for items and relationships that are of interest to humanities researchers. A SPARQL endpoint backed directly by the triplestore is one option, but presents a number of drawbacks. It can be cumbersome in SPARQL to query certain data structures that are especially useful in the humanities, and there is no standard support for permissions or for versioning of data. Queries that may return huge results also pose scalability problems. A technical solution to these problems is

proposed here. Gravsearch aims to provide the power and flexibility of a SPARQL endpoint, while providing better support for humanities data, and integrating well into a developer-friendly web-based API. Its basic design is of broad relevance, because it suggests a practical way to go beyond some limitations of the ways that humanities data has generally been made searchable.

A Gravsearch query is a virtual SPARQL query, i.e. it is processed by a server application, which translates it into one or more SPARQL queries to be processed by the triplestore. Therefore, it can offer support for data structures that are especially useful in the humanities, such as text markup and calendar-independent historical dates, and are not included in RDF standards. More generally, a Gravsearch query can use data structures that are simpler than the ones used in the triplestore, thus improving ease of use. A virtual query also allows the application to filter results according to user permissions, enforce the paging of

<sup>1</sup>The source code of the Gravsearch implementation, and its design documentation, are freely available online; see <https://www.knora.org>.

1 results to improve scalability, take into account the ver-  
2 sioning of data in the triplestore, and return responses  
3 in a form that is more convenient for web applica-  
4 tion development. The input SPARQL is independent  
5 of the triplestore implementation used, and the trans-  
6 former backend generates vendor-specific SPARQL  
7 as needed, taking into account the triplestore’s im-  
8 plementation of inference, full-text searches, and the  
9 like. Instead of simply returning a set of triples, a  
10 Gravsearch query can produce a JSON-LD response  
11 whose structure facilitates web application develop-  
12 ment.

13 The current implementation of Gravsearch is closely  
14 tied to a particular application, but its design is of more  
15 general interest to developers of RDF-based systems.  
16 It demonstrates that, for an application that manages  
17 data in an RDF triplestore and provides a web-based  
18 API, there is considerable value in a SPARQL-based  
19 API, which accepts SPARQL queries from the cli-  
20 ent, transforms them in the application, runs the trans-  
21 formed queries in the triplestore, and transforms the  
22 results before returning them to the client. It also sug-  
23 gests solutions to some of the practical issues that arise  
24 in the development of such a system. Consideration  
25 of this approach may lead others to develop similar  
26 application-specific tools, or to envisage more generic  
27 approaches to SPARQL transformation.<sup>2</sup>

### 28 1.1. Institutional and technical context

29  
30  
31 Gravsearch has been developed as part of Knora  
32 (Knowledge Organization, Representation, and An-  
33 notation), an application developed by the Data and  
34 Service Center for the Humanities (DaSCH) [6] to en-  
35 sure the long-term availability and reusability of re-  
36 search data in the humanities.<sup>3</sup> The Swiss National  
37 Science Foundation (SNSF) requires researchers to  
38 have a plan for making their research data publicly ac-  
39 cessible,<sup>4</sup> and the DaSCH was created to provide the  
40 necessary infrastructure to meet this requirement.

41 It is not feasible or cost-effective for the DaSCH to  
42 maintain a multitude of different data formats and stor-  
43 age systems over the long term. Moreover, part of the

44  
45  
46 <sup>2</sup>For some steps in this direction, see [4] and [5].

47 <sup>3</sup>The DaSCH as an institution is responsible for research data cre-  
48 ated by humanities projects. However, this does not rule out the pos-  
49 sibility that the technical solutions developed by the DaSCH are use-  
50 ful to other fields of research as well.

51 <sup>4</sup>[http://www.snf.ch/en/theSNSF/research-policies/open\\_research\\_data/Pages/default.aspx](http://www.snf.ch/en/theSNSF/research-policies/open_research_data/Pages/default.aspx)

1 DaSCH’s mission is to make all the data it stores inter-  
2 operable, so that it is possible to search for data across  
3 projects and academic disciplines in a generic way, re-  
4 gardless of the specific data structures used in each  
5 project. For example, many projects store text with  
6 markup, and it should be possible to search the markup  
7 of all these texts, regardless of whether they are letters,  
8 poems, books, or anything else. When markup con-  
9 tains references to other data, it is useful for human-  
10 ities researchers to search *through* the markup to the  
11 data it refers to, e.g. by searching for texts that refer  
12 to a person born after a certain date (as in the example  
13 in Section 2.2). It is therefore desirable to store text  
14 markup in the same way as other data, so they can be  
15 searched together.

16 The DaSCH must also mitigate the risks associated  
17 with technological and institutional change. The more  
18 its data storage is based on open standards, the easier  
19 it will be to migrate data to some other format if it  
20 becomes necessary to do so in the future.

21 With these requirements in mind, the DaSCH has  
22 chosen to store all data in an RDF triplestore as far  
23 as possible (the main exceptions being images, au-  
24 dio, and video). RDF’s flexibility allows it to accom-  
25 modate all sorts of data structures. Its standardisation,  
26 along with the variety of triplestore implementations  
27 available, helps reduce the risks of long-term preserva-  
28 tion. Even if RDF technology disappears over time, the  
29 RDF standards will make it possible to migrate RDF  
30 graphs to other types of graph storage.

31 Knora is therefore based on an RDF triplestore and  
32 a base ontology. The base ontology defines basic data  
33 types and abstract data structures; it is generic and does  
34 not make any assumptions about semantics.<sup>5</sup> Each re-  
35 search project using Knora must provide one or more  
36 ontologies defining the OWL classes and properties  
37 used in its data, and these ontologies must be derived  
38 from the Knora base ontology.

39 Knora is a server application written in Scala; it  
40 provides a web-based API that allows data to be quer-  
41 ied and updated, and supports the creation of virtual re-  
42 search environments that can work with heterogeneous  
43 research data from different disciplines.<sup>6</sup>

44 The DaSCH’s mission of long-term preservation re-  
45 quires it to minimise the risk of vendor lock-in. Knora

46  
47  
48 <sup>5</sup>For details of the Knora base ontology, see the documentation at  
49 <https://www.knora.org>.

50 <sup>6</sup>For more information on Knora, see <https://www.knora.org>. A  
51 list of projects implemented, planned, or under development using  
Knora can be found at <https://dasch.swiss/projects/>.

1 must therefore, as far as possible, avoid relying on  
2 any particular triplestore implementation, or on any  
3 type of data storage that has not been standardised and  
4 does not have multiple well-maintained implementa-  
5 tions. This means that as long as there is no widely  
6 implemented RDF standard for some of the DaSCH's  
7 requirements, such as access control and versioning,  
8 Knora has to handle these tasks itself to ensure consist-  
9 ent behaviour, regardless of which triplestore is used.<sup>7</sup>  
10 It was therefore decided, for example, to use an in-  
11 band approach to storing permissions in a role-based  
12 access control model, and to store previous versions of  
13 values as linked lists, all in standard RDF.<sup>8</sup> Over time,  
14 if new RDF standards for these features are created and  
15 widely supported, it will become possible to remove  
16 them from Knora and allow the triplestore to handle  
17 them.

## 18 1.2. Problem definition

19 Knora's API required a query language that allows  
20 for complex searches in specific projects, as well as  
21 cross-project searches using ontologies shared by dif-  
22 ferent projects, such as FOAF<sup>9</sup> or dcterms.<sup>10</sup> Re-  
23 searchers need to be able to find connections in their  
24 data across project boundaries, and developers need to  
25 present a project's data online in an interactive way to  
26 researchers and the interested public. There needs to  
27 be a flexible way to request particular graphs of data,  
28 not only for searches, but also to present combinations  
29 of resources that are known to be of interest to users.  
30 For example, the Bernoulli-Euler Online project (see  
31 Section 3) needed to present a manuscript page aligned  
32 with different layers of transcription, ranging from dip-  
33 lomatic transcriptions to translations [9]. To maximise  
34 the autonomy of web developers, Knora's API needs to  
35 make this possible without introducing project-specific  
36 API routes.

37 The query language also needs to allow queries to  
38 use data structures that are simpler than the ones stored  
39 in the triplestore. The DaSCH's requirement to have a

40 <sup>7</sup>One exception is full-text search, which is implemented by the  
41 triplestore but is not standardised.

42 <sup>8</sup>For an overview of different approaches to RDF access control,  
43 see [7]. See The Fundamentals of Semantic Versioned Querying for  
44 a proposed versioning extension to RDF. The RDF\*/SPARQL\* re-  
45 fication proposal [8] may also lead to new ways of solving these  
46 problems.

47 <sup>9</sup><http://xmlns.com/foaf/spec/>

48 <sup>10</sup>[https://www.dublincore.org/specifications/dublin-core/  
49 dcmi-terms/](https://www.dublincore.org/specifications/dublin-core/dcmi-terms/)

1 generic storage system for humanities data, as outlined  
2 in Section 1.1, introduces complexities in the the RDF  
3 data structures used in the triplestore, and these com-  
4 plexities should be hidden from clients for the sake of  
5 usability.

6 For example, in humanities research, it is useful  
7 to search for dates independently of the calendar in  
8 which they are written in source materials. If an an-  
9 cient Chinese manuscript records a solar eclipse or a  
10 sighting of a comet, and an ancient Greek manuscript  
11 records the same astronomical event at the same time,  
12 a search for texts mentioning the event and the date  
13 should find both texts, even though the date was writ-  
14 ten in different calendars. Astronomers and historians  
15 have long used a calendar-independent representation  
16 of dates, called Julian Day Numbers (JDNs), to facilit-  
17 ate such comparisons [10, 11]. A Julian Day Number  
18 is an integer, and can therefore be efficiently compared  
19 in a database query. Moreover, historical dates are of-  
20 ten imprecise, e.g. when only the year is known, but  
21 not the month or day, or when a date is known only to  
22 fall within a certain range.

23 Knora therefore stores each date in the triplestore  
24 as an instance of its `DateValue` class, which con-  
25 tains two JDNs, representing a possibly imprecise date  
26 as a date range. Each date in the range (i.e. the start  
27 date and the end date) has a precision (year, month,  
28 or day). The date range also indicates the calendar in  
29 which the date was originally recorded. This allows  
30 a SPARQL query to determine the chronological rela-  
31 tionship between two dates, regardless of the calen-  
32 dar that they were originally created in. This approach  
33 cannot handle dates whose JDN cannot be computed  
34 (e.g. dates in a calendar system whose points of refer-  
35 ence are not known); these would need to be stored in  
36 a different way. It would also be possible to extend this  
37 system to express different sorts of imprecision. How-  
38 ever, the current design is already sufficient for many  
39 historical research projects. For our purposes here, the  
40 important point is to illustrate the usefulness of hav-  
41 ing different internal and external representations of  
42 dates.<sup>11</sup>

43 <sup>11</sup>Projects such as GODOT [12] approach this task differently, by  
44 creating reference entries for calendar dates identified by URIs that  
45 can then be used in online editions to allow for searches for a spe-  
46 cific date. In our view, the approach of storing and comparing JDNs  
47 in the triplestore allows for more powerful searches, because JDNs  
48 can be compared using operators such as less than and greater than.  
49 In the future, we would like support period terms. These refer to  
50 timespans that are not expressed in terms of calendar dates, and can  
51 be represented as IRIs. For example, historians of mathematics use

JDNs are clearly not a convenient representation for clients to use. Search requests and responses should instead use whatever calendars the client prefers. The client should not have to send or receive JDNs itself, or to deal with the complexities of comparing JDN ranges in SPARQL. In the Knora API, all dates are input and output as calendar dates; Knora converts between JDNs and calendar dates using the International Components for Unicode library.<sup>12</sup>

Another example concerns text with markup, which Knora stores as RDF data (see Section 3.3). A humanities researcher might wish to search a large number of texts for, say, a particular word marked up as a noun. Knora could optimise this search by using a full-text search index. This also involves rather complex SPARQL, which is partly specific to the type of full-text indexing software being used. The client should not have to deal with these details.

SPARQL lacks other features required in this context. Knora must restrict access to data according to user permissions. It also implements a system for versioning data in the triplestore, such that the most recent version is returned by default, but the version history of resources can be requested. To improve scalability, Knora should enforce the paging of search results, rather than leaving this up to the client as in SPARQL.

### 1.3. Related work

One way of making RDF data publicly available and queryable is by means of a SPARQL endpoint backed directly by a triplestore. Prominent examples include DBpedia [16], Wikidata [17], and Europeana [3]. While this approach offers great flexibility and allows for complex queries, its drawbacks have been criticised. In a widely cited blog post [18], Dave Rogers argues that SPARQL endpoints are an inherently poor design that cannot possibly scale, and that RESTful APIs should be used instead. For example, a SPARQL endpoint allows a client to request all the data in the repository; this could easily place unreasonable demands on the server, particularly if many such requests are submitted concurrently.

---

the expression ‘Erste Petersburger Periode’ [13] to refer to the first timespan Leonhard Euler lived in St. Petersburg. Projects such as ChronOntology [14] or PeriodO [15] relate period terms to timespans and geographical regions. Period terms can also be made to correspond to calendar dates stored as JDNs, as long as the calendar date representation can express the imprecision required in each term.

<sup>12</sup><http://site.icu-project.org/home>

GraphQL [19] is a newer development and – despite its name – not restricted to graph databases. It is meant to be a query language that integrates different API endpoints. Instead of making several requests to different APIs and processing the results individually, GraphQL is intended to allow the client to make a single request that defines the structure of the expected response. (See section 2.1 for a discussion of extensions to GraphQL for querying RDF.)

From our perspective, triplestore-backed SPARQL endpoints and GraphQL both have limitations that make them unsuitable for Knora and for humanities data in general. They assume that the data structures in the triplestore are the same as the ones to be returned to the client. They offer no standard way to restrict query results according to the client’s permissions. They do not enforce the paging of results, but leave this to the client. And they provide no way to work with data that has a version history (so that ordinary queries return only the latest version of each item). These requirements led us to develop a different approach.

## 2. A hybrid between a SPARQL endpoint and a web API

One option would be to create a domain-specific language, but it was simpler to use SPARQL, leveraging its standardisation and library support, while integrating it into Knora’s web API. Gravsearch therefore accepts as input a subset of standard SPARQL query syntax, and requires queries to follow certain additional rules.

Gravsearch is thus a hybrid between a SPARQL endpoint and a web API, aimed at combining the advantages of both. By supporting SPARQL syntax, it enables clients to submit queries based on complex graph patterns given in a `WHERE` clause. By supporting SPARQL `CONSTRUCT` queries, it allows each query to specify a complex graph structure to be returned in the query results. At the same time, the application is able to add additional functionality not supported in standard SPARQL. The Knora API server processes the input query, transforming it into one or more SPARQL queries that are executed by the triplestore, using data structures that are more complex than the ones in the input query. It then processes and transforms the triples returned by the triplestore, this time converting complex data structures into simpler ones, to construct the response.

1 This extra layer of processing enables Gravsearch to  
2 avoid the disadvantages of SPARQL endpoints backed  
3 directly by a triplestore, and to provide additional  
4 features. Certain data structures can be queried in a  
5 more convenient way, results are filtered according to  
6 the user's permissions, the versioning of data in the  
7 triplestore is taken into account (only the most recent  
8 version of the data is returned), and scalability is im-  
9 proved by returning results in pages of limited size.

### 11 2.1. Scope of Gravsearch

13 Gravsearch was developed as part of Knora and is  
14 used as part of the Knora API. Knora is an integrated  
15 system, taking care of creating, updating and reading  
16 data by mediating between the triplestore and the cli-  
17 ent. Gravsearch is not meant to replace direct commu-  
18 nication with a triplestore via a SPARQL endpoint, but  
19 rather to provide an integrated system with a flexible  
20 way of querying data. The result of a Gravsearch query  
21 is returned in the same format used by the rest of the  
22 Knora API, making it suitable for web applications.

23 HyperGraphQL [20], an extension to GraphQL,  
24 makes it possible to query SPARQL endpoints using  
25 GraphQL queries, by converting them to SPARQL. Its  
26 intended advantages include the reduction of complex-  
27 ity on the client side and a more controlled way of  
28 accessing a SPARQL endpoint, avoiding some of the  
29 problems discussed in Rogers's blog post [21]. A sim-  
30 ilar approach has been taken with GraphQL-LD [22],  
31 differing from HyperGraphQL in that GraphQL quer-  
32 ies are translated to SPARQL on the client-side. We  
33 think that both HyperGraphQL and GraphQL-LD  
34 could be viable approaches to generating SPARQL that  
35 is further processed before being sent to the triplestore,  
36 thus integrating with Gravsearch. We consider this a  
37 promising avenue for future development.

38 When the full flexibility of SPARQL is needed, pro-  
39 jects always have the possibility to make their data  
40 openly accessible via a SPARQL endpoint. In such a  
41 case, a named graph can be made for queries but not  
42 for updates, and possibly transformed into a simpler  
43 data model than used internally in order to facilitate  
44 integration with other data sets.

45 It is true that, even with pagination, it is possible  
46 to write 'inefficient or complex SPARQL that returns  
47 only a few results' [18], and this could also be true of  
48 Gravsearch queries. As a last resort, some triplestores  
49 provide a way to set arbitrary limits on execution time  
50 or the number of triples or rows returned, but this still  
51 means consuming significant resources before reject-

1 ing the query. Moreover, setting a limit on the number  
2 of triples returned by a CONSTRUCT query can cause  
3 the triplestore to return incomplete entities, with no in-  
4 dication to the client that the response has been trun-  
5 cated.

6 It would be better to reject a badly written query  
7 without running it, and users would appreciate error  
8 messages explaining what needs to be changed to im-  
9 prove the query. The presence of an application layer  
10 that analyses and transforms the input query provides  
11 opportunities to do this. This is why, for example,  
12 Gravsearch currently does not allow subqueries. It  
13 would also be possible to reject a triple pattern like `?s  
14 ?p ?o` whose variables are not restricted by any other  
15 pattern. We see this kind of analysis as a promising  
16 topic for future research.

17 Although SELECT queries are not currently support-  
18 ed, if tabular output is desired, (e.g. for statistical anal-  
19 ysis), the results of a CONSTRUCT query can be con-  
20 verted into a table by combining results pages and con-  
21 verting the RDF output to tabular form. This could be  
22 done either in the client or on the server.

### 23 2.2. Ontology schemas

24 A design goal of Gravsearch is to enable queries to  
25 work with data structures that are simpler than the ones  
26 actually used in the triplestore, thus hiding some com-  
27 plexity from the user. To make this possible, Knora  
28 implements *ontology schemas*. Each ontology schema  
29 provides a different view on ontologies and data that  
30 exist in the triplestore. The term *internal schema* refers  
31 to the structures that are actually in the triplestore, and  
32 *external schema* refers to a view that transforms these  
33 structures in some way for use in a web API.

34 Knora's built-in ontologies as well as all project-  
35 specific ontologies are available in each schema, but  
36 in the triplestore, all ontologies are stored only in the  
37 internal schema; the other schemas are generated on  
38 the fly as needed. An ontology always has a differ-  
39 ent IRI in each schema, but the triplestore sees only  
40 the IRIs for the internal schema. In Knora, the IRIs of  
41 built-in and project-specific ontologies must conform  
42 to certain patterns; this allows Knora to convert onto-  
43 logic IRIs automatically, following simple rules, when  
44 converting ontologies and data between schemas. An  
45 additional convention is that Knora's base ontology  
46 is called `knora-base` in the internal schema, and  
47 `knora-api` in the external schemas.

48 In the internal schema, the smallest unit of research  
49 data is a Knora `knora-base:Value`, which is an  
50  
51

RDF entity that has an IRI. Subclasses of `Value` represent different types of data (text, integers, dates, and so on). If a client wishes to update a value via the Knora API, it needs to know the value's IRI. However, a Knora `Value` also contains information that is represented in a way that is not convenient for clients to manipulate (e.g. dates are stored as JDNs, as mentioned above). Therefore, Knora provides an external schema called the `knora-api complex schema`, in which each value has an IRI, but its contents are represented in more convenient form (e.g. calendar dates are used instead of JDNs).

For clients that need a read-only view of the data, Knora provides a *simple schema*, in which there is no `Value` class; instead, Knora values are represented as literal datatypes such as `xsd:string`, and the metadata attached to each value is hidden. An advantage of the simple schema is that it facilitates the use of standard ontologies such as `dcterms`, in which values are typically represented as datatypes without their own metadata. For example, if a property is defined in Knora as a subproperty of `dcterms:title`, its object in Knora will internally be a `Knora TextValue` with attached metadata, but a Gravsearch query in the simple schema can treat it as a literal, in keeping with its definition in `dcterms`.

Gravsearch queries can thus be written in either of Knora's external schemas, and results can also be returned in either of these schemas.

To illustrate the differences between these schemas, Appendix A shows a Gravsearch query that searches for a letter in the Bernoulli-Euler Online (BEOL) project (see Section 3), using the complex schema. The query searches for a letter from the mathematician Anders Johan Lexell (identified by his IRI), specifying that the text of the letter must refer to a person whose birthdate is after 1706 CE. To do this, the query in Listing 1 searches the triples representing the text's standoff markup (see Section 3.3) for a link to a `beol:person` resource whose birthdate is greater than that date.

To run this query, the Gravsearch transformer generates a SPARQL `CONSTRUCT` query; part of the triplestore's response is shown in Listing 2 (the full response contains many more details such as permissions, timestamps, and so on). A `beol:person` representing Leonhard Euler is returned, because he is identified in the text markup and was born after 1706. Note that the date values are represented as ranges of JDNs, and each markup tag is represented as an RDF entity.

If the client requested a JSON-LD response in the complex schema, it will look like Listing 3. Here each date is represented as a calendar date, the text with its markup has been converted to an XML document (one of several possible representations that the client can request), and each link to a `beol:person` resource is represented as a JSON-LD object containing the target resource as a nested object. Each resource is accompanied by additional metadata added by the Knora API server (some of which is not shown here), e.g. an ARK URL serving as a permanent link to the resource.

This JSON-LD representation in the complex schema is designed to be convenient to use in the development of web applications. The ARK URLs in this listing can be opened in a web browser to show an example of such an application.<sup>13</sup>

If the client requested a JSON-LD response in the simple schema, it will look like Listing 4. (The client can also request an equivalent response in Turtle or RDF/XML.) Here, the structure of the response has been simplified, so that values are represented as simple literals. For example, the object of `beol:title` is a string literal rather than a `TextValue`. Moreover, if the client dereferences the BEOL ontology IRI given in the JSON-LD context,<sup>14</sup> the definition of `beol:title` says that it is a datatype property, and that it is a subproperty of `dcterms:title`. A client that knows about `dcterms:title`, and processes titles in some particular way, can then apply the same processing to the `beol:title`. As this example shows, the simple schema lends itself to use in applications that rely mainly on standard ontologies to integrate data from different sources.

The simple schema also makes it possible to use standard ontologies in the Gravsearch query itself. Listing 5 shows a Gravsearch query that uses the FOAF ontology to search the BEOL project data for `foaf:Person` objects whose `foaf:familyName` is 'Euler', without using the BEOL ontology at all. This works because `beol:person` is a subclass of `foaf:Person`, and `beol:hasFamilyName` is a subproperty of `foaf:familyName`. Depending on the triplestore's inference capabilities, Knora either uses the triplestore's RDFS inference or expands the query itself using property path syntax.

<sup>13</sup>For example, see <http://ark.dasch.swiss/ark:/72163/1/0801/25Ro4c5gSIioLstzJ25FiAC> for a representation of the letter.

<sup>14</sup><http://api.dasch.swiss/ontology/0801/beol/simple/v2>

The query in Listing 5 specifies that the object of `foaf:familyName` is an `xsd:string`. This is consistent with the definition of the subproperty `beol:hasFamilyName` in the simple schema. In the internal schema, the object of that property is the IRI of a `Knora TextValue`, which contains a string as well as permissions and other metadata. When Gravsearch generates SPARQL from the input query, it transforms the `WHERE` clause of the input query to handle this difference (see Section 2.6). This design allows standard ontologies using simple literal values, such as FOAF, to be used in Gravsearch queries, even though Knora actually stores values as IRIs with attached metadata.

In short, Gravsearch transforms queries, ontologies, and data on the fly according to the ontology schema that the client is using. The implementation is partly object-oriented (the Scala classes representing different sorts of RDF entities have methods for generating representations of themselves in different schemas) and partly based on sets of transformation rules for each schema (e.g. to remove certain properties and add others, and to change OWL cardinalities). Additional external schemas could be added in the future.

### 2.3. Permissions

In Knora, each resource and each value has role-based permissions attached to it. Internally, permissions are represented as string literals in a compact format that optimises query performance. For example, a Knora value could contain this triple:

```
<http://rdfh.ch/0001/R5qJ6oPZPV>
  knora-base:hasPermissions
    "V http://rdfh.ch/groups/00FF/reviewer" .
```

This means that the value can be viewed by members of the specified group. With a SPARQL endpoint backed directly by a triplestore, there would be no way to prevent other users from querying the value. Therefore, the application must filter query results according to user permissions.

To determine whether a particular user can view the value, Knora must compute the intersection of the set of groups that the user belongs to and the set of groups that have view permission on the value. If not, Knora removes the value (and the resource that contains it) from the results of the Gravsearch query.

### 2.4. Versioning

Internally, a resource is connected only to the current version of each of its values. Each value version is connected to the previous version via the property `previousVersion`, so that the versions form a linked list. When a client requests a single resource with its values via the Knora API, the client can specify a version timestamp. Knora then generates a SPARQL query that traverses the linked list to retrieve the values that the resource had at the specified time.

Gravsearch is designed to query only current data. This is easily achieved, because the only way to obtain a value in Gravsearch is to follow the connection between the resource and the value, which is always the current version. Knora's external ontology schemas do not expose the version history data structure at all (e.g. they do not provide the property `previousVersion`). Therefore, the client cannot use `previousVersion` to query a past version of a value, which would be possible if they could submit SPARQL directly to the triplestore. A specific (non-Gravsearch) API request enables clients to access the version history of a value. The version history data structure is hidden from clients, enabling it to be changed in the future as RDF technology develops.

### 2.5. Gravsearch syntax and semantics

Syntactically, a Gravsearch query is a SPARQL `CONSTRUCT` query. Thus it supports arbitrarily complex search criteria. One could, for example, search for persons whose works have been published by a publisher that is located in a particular city. A `CONSTRUCT` query also allows the client to specify, for each resource that matches the search criteria, which values of the resource should be returned in the search results.

Results are returned by default as a JSON-LD array, with one element per search result. Each search result contains the 'main' or top-level resource that matched the query. If the query requests other resources that are connected to the main resource, these are nested as JSON-LD objects within the main resource. To make this possible, a Gravsearch query must specify (in the `CONSTRUCT` clause) which variable refers to the main resource. The resulting tree structure is generally more useful to web application clients than the flat set of RDF statements returned by SPARQL endpoints.

Gravsearch uses the SPARQL constructs `ORDER BY` and `OFFSET` to enable the client to step through

pages of search results, and does not allow `LIMIT`. The client can use `ORDER BY` with one or more variables to determine the order in which results will be returned, and `OFFSET` to specify which page of results should be returned. (This is different from the standard SPARQL `OFFSET`, which counts solutions to the pattern in the `WHERE` clause, rather than pages of results.) The maximum number of results per page is configurable in the application's settings, and cannot be controlled by the client. The server indicates that more pages of results may be available by including a boolean property `mayHaveMoreResults` in the response. To retrieve all results, the client continues requesting pages until this property is absent (false); it is possible that the last page will be empty.

## 2.6. Processing and execution of a Gravsearch query

In processing Gravsearch queries, the API server is free to use a SPARQL design that best suits the performance characteristics of the triplestore. For example, as described below, our implementation transforms each input query into multiple SPARQL queries that are run in the triplestore, and generates different SPARQL for different triplestores. Clients and users need not be aware of this.

In theory, such transformations could be implemented using a more generic rule system, such as Rule Interchange Format (RIF),<sup>15</sup> which is intended to be used for this purpose by the GeoSPARQL standard.<sup>16</sup> However, RIF is not widely implemented, and even most GeoSPARQL implementations do not support query rewriting [23].<sup>17</sup>

Rather than introduce an additional programming language for query transformation, we chose to implement a traditional compiler design in Scala, taking advantage of the RDF4J<sup>18</sup> SPARQL parser and Knora's extensive built-in support for working with OWL ontologies. An abstract syntax tree (AST) is constructed, type information is collected from it (see Section 2.6.3), and it is passed through a sequence of transformations, each of which recursively traverses the AST, transforming graph patterns to produce a

<sup>15</sup><https://www.w3.org/TR/rif-overview/>

<sup>16</sup><https://www.opengeospatial.org/standards/geosparql>

<sup>17</sup>Another example of a mapping language for RDF is R2RML, which intended for defining mappings between relational databases and RDF datasets. Similarly, EPNet [24] maps SPARQL queries onto heterogeneous queries in different sorts of databases, and Ontop [25] allows for querying relational databases with SPARQL.

<sup>18</sup><https://rdf4j.org>

new AST. There are different backends for different triplestores, which take into account triplestore-specific features such as RDFS inference capabilities and full-text search. We are currently exploring adding an optimisation phase to reorder patterns in the `WHERE` clause in cases where the triplestore's query optimiser may not do so effectively.

For the reasons explained in Section 1.2, the generated SPARQL is considerably more complex than the provided Gravsearch query, and deals with data structures in the internal schema. Each Gravsearch query is converted to two SPARQL queries to improve performance. First, a `SELECT` query (prequery) is generated, to identify a page of matching resources. Then a `CONSTRUCT` query (main query) is generated, to retrieve the requested values of those resources. Figure 1 illustrates this process graphically.

### 2.6.1. Generation of the Prequery

The prequery's purpose is to get one page of IRIs of matching resources and values. It is a `SELECT` query whose `WHERE` clause is generated by transforming the statements of the input query's `WHERE` clause. Thus it contains all the restrictions of the input query, transformed into the internal schema (see 2.2), as well as additional statements, e.g. to implement RDFS inference if the triplestore's inference capabilities are not adequate, and to ignore resources and values that have been marked as deleted. The prequery's `SELECT` clause is generated automatically.

The prequery's result consists of the IRIs of one page of matching main resources, along with their values and linked resources. Since correct paging requires the query to return one row per matching main resource, the results are grouped by main resource IRI, and the IRIs of matching values and of linked resources are aggregated using `GROUP_CONCAT`. The results are ordered by the criteria in the `ORDER BY` clause of the input query, as well as by main resource IRI (to ensure deterministic results).

The following sample Gravsearch query gets all the entries with sequence number equal to 10 from different manuscripts.

```

PREFIX beol:
  <http://beol.dasch.swiss/ontology/
    0801/beol/simple/v2#>
PREFIX knora-api:
  <http://api.knora.org/ontology/
    knora-api/simple/v2#>

CONSTRUCT {

```



```

1   ?entry knora-api:isMainResource true .      CONSTRUCT {
2   ?entry beol:manuscriptEntryOf ?manuscript . ...
3   ?entry beol:seqnum ?seqnum .                }
4 } WHERE {                                     WHERE {
5   ?entry a beol:manuscriptEntry .             {
6   ?entry beol:manuscriptEntryOf               VALUES ?mainResourceVar {
7     ?manuscript .                             <http://rdfh.ch/0803/1>
8   ?manuscript a beol:manuscript .             <http://rdfh.ch/0803/2> ...
9   ?entry beol:seqnum ?seqnum .                }
10  FILTER(?seqnum = 10)                       ?mainResourceVar a knora-base:Resource .
11 }                                              ...
12 }                                              } UNION { ...

```

The resulting prequery looks like this:

```

13 SELECT DISTINCT
14   ?entry
15   (GROUP_CONCAT(DISTINCT(?manuscript);
16     SEPARATOR='\u001F')
17   AS ?manuscript__Concat)
18   (GROUP_CONCAT(DISTINCT(?seqnum);
19     SEPARATOR='\u001F')
20   AS ?seqnum__Concat)
21   (GROUP_CONCAT(DISTINCT(
22     ?manuscriptEntryOf__LinkValue);
23     SEPARATOR='\u001F')
24   AS
25     ?manuscriptEntryOf__LinkValue__Concat)
26 WHERE {...}
27 GROUP BY ?entry
28 ORDER BY ASC(?entry)
29 LIMIT 25

```

The variables in the prequery's SELECT clause represent the IRIs of the matching main resources and the linked resources and values present in the input query's WHERE clause. Besides the IRIs of the main resources, all IRIs are returned concatenated as the result of an aggregation function.

### 2.6.2. Generation of the Main Query

The result of the prequery is a collection of IRIs of matching resources and values, grouped by main resource. The main query is then generated; it is a SPARQL CONSTRUCT query that specifically requests the resources and values identified by those IRIs.

Unlike the prequery, the main query's WHERE clause is not based on the input query; its structure is always the same. The application creates statements using the VALUES keyword, and inserts the IRIs returned by the prequery for different categories of information, using UNION blocks: the main resources, linked resources, values, and standoff markup (if necessary).

The code below shows a snippet from the main query.

The main query's results are the contents of a page of resources and values that matched the input query's WHERE clause. The application can now perform permission checks and filter out those resources and values that the client does not have permission to see.

Finally, the application orders the main query's results according to the order in which the main resources were returned by the prequery, returning a JSON-LD array with one element per main resource.

### 2.6.3. Type checking and inference

SPARQL does not provide type checking; if a SPARQL query uses a property with an object that is not compatible with the property definition, the query will simply return no results.

However, Gravsearch requires the types of the entities used in a query so it can generate the correct SPARQL. Specifically, if a query uses the simple schema, it needs to be expanded to work with the internal schema, by taking into account an additional layer of value entities rather than simple literal values. The compiler therefore needs to know:

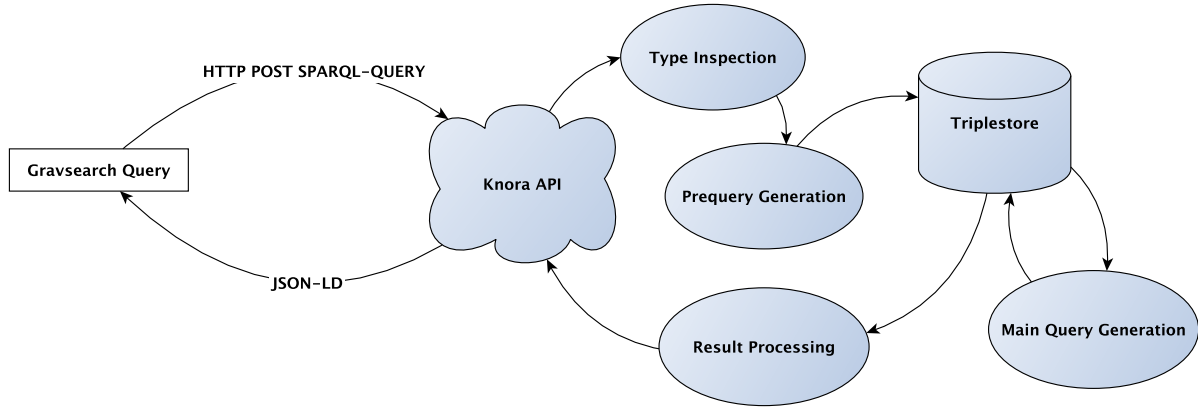
- The type of each variable or entity IRI used as the subject or object of a statement.
- The type that is expected as the object of each property used.

For the sake of efficiency, it is desirable to obtain this information without doing additional SPARQL queries, using only the information provided in the query itself along with the available ontologies in the triplestore (which Knora keeps in memory).<sup>19</sup>

Gravsearch therefore implements a simple type inference algorithm, focusing on identifying the types

<sup>19</sup>This approach contrasts with a mechanism such as SHACL [26], which can run SPARQL queries in the triplestore.

Figure 1. Processing and Execution of a Gravsearch Query



that are relevant to the compiler.<sup>20</sup> It first collects the set of all entities whose types need to be determined (*typeable entities*) in the WHERE clause of the input query. It then runs the WHERE clause through a pipeline of *type inspectors*. Each inspector implements a particular mechanism for identifying types, and passes an intermediate result to the next inspector in the pipeline.

At the end of the pipeline, each typeable entity should have exactly one type from the set of types that are relevant to the compiler. Unlike a SPARQL endpoint, if Gravsearch cannot determine the type of an entity, or finds that an entity has been used inconsistently (i.e. with two different types in that set), it returns an error message rather than an empty response.

The first type inspector reads type annotations. Two annotations are supported:

- `rdf:type`, used to specify the types of entities that are used as subjects or objects.
- `knora-api:objectType`, which can be used to specify the expected type of the object of a non-Knora property such as `dcterms:title`.

Annotations are needed when a query uses a non-Knora ontology such as FOAF or `dcterms`. In the future, it would be useful for Knora to take type information from standard non-Knora ontologies, making these annotations unnecessary.

The second inspector in the pipeline infers types by using class and property definitions in ontologies. It

runs each typeable entity through a pipeline of inference rules. These include rules such as the following:

- The type of a property’s object is inferred from the expected object type of the property (which is specified in the definition of each Knora property).
- The expected object type of a property is inferred from the type of its actual object. (Thus if the query specifies `?book dcterms:title ?title` and `?title a xsd:string`, the rule infers that `dcterms:title` expects an `xsd:string` object.)
- If a FILTER expression compares two entities, they are inferred to have the same type.
- Function arguments are inferred to have the required types for the function.

Since the output of one rule may allow another rule to infer additional information, the pipeline of inference rules is run repeatedly until no new information can be determined. In practice, two iterations are sufficient for most realistic queries.

Appendix B shows an input query in the simple schema. It searches for books that have a particular publisher (identified by IRI), and returns them along with the family names of all the persons that have some connection with those books (e.g. as author or editor).

In this example, the definition of the property `hasPublisher` specifies that its object must be a `Publisher`, allowing Gravsearch to infer the type of the resource identified by the specified IRI. Similarly, the definition of `hasFamilyName` specifies that its subject must be a `Person` and its object must be a `Knora TextValue`; this allows the types of `?person` and `?familyName` to be inferred. Once

<sup>20</sup>The Knora base ontology determines the set of types that the algorithm needs to identify, and thus simplifies the algorithm. For an attempt at more complete type inference for SPARQL queries, see [27].

1 the type of `?person` is known, the object type of  
2 `?linkProp` can be inferred.

### 3. Use case from the Bernoulli-Euler Online project

7 One project that is using Gravsearch is Bernoulli-Euler Online (BEOL),<sup>21</sup> a digital edition project focusing on 17th- and 18th-century primary sources in mathematics. BEOL integrates written sources relating to members of the Bernoulli dynasty and Leonhard Euler into a web application based on Knora, with data stored in an RDF triplestore. The BEOL web site provides a user interface that enables users to search and view these texts in a variety of ways. It offers a menu of common queries that internally generate Gravsearch using templates, and the user can also build a custom query using a graphical search interface, which also generates Gravsearch internally.

#### 3.1. Example 1: finding correspondence between two mathematicians

25 Most of the texts that are currently integrated in the BEOL platform are letters exchanged between mathematicians. On the project's landing page, we would like to present the letters arranged by their authors and recipients. With Gravsearch, it is not necessary to make a custom API operation for this kind of query in Knora. Instead, a Gravsearch template can be used, with variables for the correspondents.

33 Appendix C shows a template for a Gravsearch query that finds all the letters exchanged between two persons. Each person is represented as a resource in the triplestore. It would be possible to use the IRIs of these resources to identify mathematicians, but since these IRIs are not yet stable during development, it is more convenient to use the property `beol:hasIAFIdentifier`, whose value is an Integrated Authority File (IAF) identifier (maintained by the German National Library), a number that uniquely identifies that person. This example thus illustrates searching for resources that have links to other resources that have certain properties. The user chooses the names of two mathematicians from a menu in a web browser, and the user interface then processes the template, substituting the IAF identifiers of those two mathematicians for the placeholders `#{iaf1}` and

1 `#{iaf2}`. The result of processing the template is a Gravsearch query, which the user interface submits to the Knora API server. This query specifies that the author and recipient of each matching letter must have one of those two IAF identifiers. The results are sorted by date. The page number `#{offset}` is initially set to 0; as the user scrolls, the page number is incremented and the query is run again to load more results.

9 This query is simple enough to be written in the simple schema. For example, this allows the object of `beol:hasIAFIdentifier` to be treated as a string literal. Internally, this is an object property. Its object is an entity belonging to the class `knora-base:TextValue`, and has predicates and objects of its own. This extra level of complexity is hidden from the client in the simple schema. After we substitute the IAF identifiers of Leonhard Euler and Christian Goldbach for the placeholders in the template, the input query contains:

```
21 ?author beol:hasIAFIdentifier
22 ?authorIAF .
23 FILTER(?authorIAF =
24 " (DE-588)118531379" ||
25 ?authorIAF = " (DE-588)118696149")
```

26 Gravsearch transforms these two lines to the following SPARQL:

```
28 ?author beol:hasIAFIdentifier ?authorIAF .
29 ?authorIAF knora-base:isDeleted false .
30 ?authorIAF knora-base:valueHasString
31 ?authorIAF__valueHasString .
32 FILTER(?authorIAF__valueHasString =
33 " (DE-588)118531379"^^xsd:string ||
34 ?authorIAF__valueHasString =
35 " (DE-588)118696149"^^xsd:string)
```

37 Since values in Knora can be marked as deleted, the generated query uses `knora-base:isDeleted false` to exclude deleted values. It then uses the generated variable `?authorIAF__valueHasString` to match the content of the `TextValue`.

#### 3.2. Example 2: a user interface for creating queries

42 Users can also create custom queries that are not based on a predefined template. For this purpose, a user-interface widget generates Gravsearch, without requiring the user to write any code (Appendix E).

43 For example, a user can create a query to search for all letters written since 1 January 1700 CE (the user specifies the Gregorian calendar) by Johann I

51 <sup>21</sup><https://beol.dasch.swiss/>

1 Bernoulli, that mention Leonhard Euler but not Daniel  
2 I Bernoulli, and that contain the word *Geometria*. The  
3 user can choose to order the results by date. The  
4 web-based user interface generates a Gravsearch query  
5 based on the search criteria (Appendix D).

6 In generating SPARQL to perform the requested  
7 search, the Gravsearch compiler converts the date  
8 comparison to one that uses a JDN. In the example,  
9 the input SPARQL requests a date greater than a date  
10 literal in the Gregorian calendar:

```
11 FILTER(?date >=  
12 "GREGORIAN:1700-1-1"^^knora-api:Date)
```

14 Gravsearch converts this to a JDN comparison. The  
15 Gregorian date 1 January 1700 is converted to the JDN  
16 2341973. In the generated SPARQL, a matching date's  
17 end point must be greater than or equal to that JDN:

```
19 ?date knora-base:valueHasEndJDN  
20 ?date__valueHasEndJDN .  
21 FILTER(?date__valueHasEndJDN >=  
22 "2341973"^^xsd:integer)
```

24 Other date comparisons work as follows:

- 25 – Two `DateValue` objects are considered equal if  
26 there is any overlap between their date ranges.
- 27 – Two `DateValue` objects are considered unequal  
28 if there is no overlap between their date ranges.
- 29 – `DateValue A` is considered to be less than  
30 `DateValue B` if A's end date is less than B's  
31 start date.

33 To specify that the text of the letter must contain  
34 the word *Geometria*, the input SPARQL uses the func-  
35 tion `knora-api:matchText`, which is provided  
36 by Gravsearch:

```
37 FILTER knora-api:matchText(?text,  
38 "Geometria")
```

40 Gravsearch converts this function to triplestore-  
41 specific SPARQL that (unlike the standard SPARQL  
42 `CONTAINS` function) performs the query using a full-  
43 text search index. For example, with the GraphDB  
44 triplestore using the Lucene full-text indexer, the gen-  
45 erated query contains:

```
47 ?text knora-base:valueHasString  
48 ?text__valueHasString .  
49 ?text__valueHasString  
50 lucene:fullTextSearchIndex  
51 "Geometria"^^xsd:string .
```

### 3.3. Example 3: Searching for text markup

3 Here we are looking for a text containing the word  
4 *Acta* that is marked up as a bibliographical reference.<sup>22</sup>

5 Knora stores text markup as 'standoff markup': each  
6 markup tag is represented as an entity in the triplestore,  
7 with start and end positions referring to a substring  
8 in the text. This makes it straightforward to represent  
9 non-hierarchical structures in markup,<sup>23</sup> and makes it  
10 possible for queries to combine criteria referring to  
11 text markup with criteria referring to other entities in  
12 the triplestore, including links within text markup that  
13 point to RDF resources outside the text. Projects may  
14 define own standoff entities in their project-specific  
15 ontologies, deriving them from the types defined by  
16 the Knora base ontology.

17 To search for text markup, the input query must be  
18 written in the complex schema. The input query uses  
19 the `matchTextInStandoff` function provided by  
20 Gravsearch:

```
22 ?text knora-api:textValueHasStandoff  
23 ?standoffBibliographyTag .  
24 ?standoffBibliographyTag a  
25 beol:StandoffBibliographyTag .  
26 FILTER knora-api:matchTextInStandoff(  
27 ?text,  
28 ?standoffBibliographyTag,  
29 "Acta")
```

30 Gravsearch translates this `FILTER` into two opera-  
31 tions:

- 32 1. An optimisation that searches in the full-text  
33 search index to find all texts containing this  
34 word.
- 35 2. A regular expression match that determines  
36 whether, in each text, the word is located within  
37 a substring that is marked up as a paragraph.

38 The resulting generated SPARQL looks like this<sup>24</sup>:

```
39 ?text knora-base:valueHasString  
40 ?text__valueHasString .  
41 ?text__valueHasString  
42 lucene:fullTextSearchIndex
```

46 <sup>22</sup>'Acta' refers to an article published by Jacob Bernoulli in *Acta*  
47 *Eruditorum* in December 1695.

48 <sup>23</sup>See the TEI guidelines [28, chapter 20, Non-hierarchical Struc-  
49 tures] for a discussion of this problem.

50 <sup>24</sup>Knora uses 0-based indexes in standoff markup, but SPARQL  
51 uses 1-based indexes: <https://www.w3.org/TR/xpath-functions/#func-substring>.

```

1      "Acta"^^xsd:string .
2      ?standoffTag
3      knora-base:standoffTagHasStart
4      ?standoffTag__start .
5      ?standoffTag
6      knora-base:standoffTagHasEnd
7      ?standoffTag__end .
8      BIND (substr(?text__valueHasString,
9      ?standoffTag__start + 1,
10     ?standoffTag__end -
11     ?standoffTag__start)
12     AS ?standoffTag__markedUp)
13     FILTER(regex(?standoffTag__markedUp,
14     "Acta", "i"))
15

```

#### 4. Conclusion

To ensure the long-term accessibility of research data, a case can be made for storing nearly all data (with the exception of a few binary file formats) in RDF, in a way that works with any standards-compliant triplestore, and avoiding non-standard technologies and vendor lock-in. However, this approach implies that clients cannot be given direct access to the triplestore, both because the necessary generic data structures are inconvenient for clients to use, and because features such as versioning and access control cannot be handled by the triplestore in a standard way. There are also scalability issues associated with using a SPARQL endpoint backed directly by the triplestore. Yet humanities researchers want to be able to do the sorts of powerful graph searches that they can do in SPARQL.

The solution proposed here is to let users submit SPARQL, but to a virtual endpoint (the Knora API) rather than to the triplestore. In this way, the Knora API can serve its main purpose of keeping research data available for the long term without relying on vendor-specific triplestore features, while providing flexible, controllable, and consistent access to data. By sending SPARQL queries to the Knora API, users can access only the data they have permission to see, only the current version of the data is served, and implementation details are hidden. Results are returned using a paging mechanism controlled by the Knora API. By default, the query results are returned in a tree structure in JSON-LD. This data structure is suitable for web application development, while maintaining machine-readability of the data as well as interoperability with other RDF-based tools.

#### Acknowledgements

This work was supported by the Swiss National Science Foundation (166072) and the Swiss Data and Service Center for the Humanities.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51

## Appendix A. Ontology schemas

```

1 PREFIX knora-api: <http://api.knora.org/ontology/knora-api/v2#>
2 PREFIX knora-api-simple: <http://api.knora.org/ontology/knora-api/simple/v2#>
3 PREFIX beol: <http://api.dasch.swiss/ontology/0801/beol/v2#>
4
5 CONSTRUCT {
6   ?letter knora-api:isMainResource true .
7   ?letter beol:title ?title .
8   ?letter beol:hasAuthor <http://rdfh.ch/0801/1PkbSo2nRIeOK9ISp6JSdg> .
9   ?letter beol:hasRecipient ?recipient .
10  ?letter beol:creationDate ?creationDate .
11  ?letter beol:hasText ?text .
12  ?letter knora-api:hasStandoffLinkTo ?person .
13  ?person beol:hasBirthDate ?birthDate .
14 } WHERE {
15   ?letter a beol:letter .
16   ?letter beol:title ?title .
17   ?letter beol:hasAuthor <http://rdfh.ch/0801/1PkbSo2nRIeOK9ISp6JSdg> .
18   ?letter beol:hasRecipient ?recipient .
19   ?letter beol:creationDate ?creationDate .
20   ?letter beol:hasText ?text .
21   ?text knora-api:textValueHasStandoff ?entityTag .
22   ?entityTag a beol:StandoffEntityTag .
23   FILTER knora-api:standoffLink(?letter, ?entityTag, ?person)
24   ?person a beol:person .
25   ?person beol:hasBirthDate ?birthDate .
26   FILTER(knora-api:toSimpleDate(?birthDate) >= "GREGORIAN:1706 CE"^^knora-api-simple:Date)
27   ?letter knora-api:hasStandoffLinkTo ?person .
28 }

```

Listing 1: Searching for a letter whose text mentions a person born after 1706

```

31 @prefix knora-base: <http://www.knora.org/ontology/knora-base#> .
32 @prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
33 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
34 @prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
35 @prefix beol: <http://www.knora.org/ontology/0801/beol#>
36
37 <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA> a beol:letter;
38   rdfs:label "Lexell à Condorcet";
39   beol:title <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/LWf9YstAQ4iq4CI1MmOzUA>;
40   beol:creationDate <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/M4eNIJ1MR_eB9mvTGtaZng>;
41   beol:hasAuthorValue <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/7ExuBcDbQZyRJRbAcORyKg>;
42   beol:hasRecipientValue <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/wCaspGWvS6SZXeuZy3h9MA>;
43   beol:hasText <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/im2jaDG_RKcX999ozN3RZg> .
44
45 <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/LWf9YstAQ4iq4CI1MmOzUA> a knora-base:TextValue;
46   knora-base:valueHasString "Annexe 1: Lexell à Condorcet, Saint-Pétersbourg, 2 (13) décembre 1775" .
47
48 <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/M4eNIJ1MR_eB9mvTGtaZng> a knora-base:DateValue;
49   knora-base:valueHasCalendar "JULIAN";
50   knora-base:valueHasEndJDN 2369712;
51   knora-base:valueHasEndPrecision "DAY";
52   knora-base:valueHasStartJDN 2369712;
53   knora-base:valueHasStartPrecision "DAY" .
54
55 <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/7ExuBcDbQZyRJRbAcORyKg> a knora-base:LinkValue;

```

```

1   rdf:object <http://rdfh.ch/0801/1PkbSo2nRIeOK9ISp6JSdg>;
2   rdf:predicate beol:hasAuthor;
3   rdf:subject <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA> .
4
5   <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/wCaspGWvS6SZXeuZy3h9MA> a knora-base:LinkValue;
6   rdf:object <http://rdfh.ch/0801/7z7sHo5aTLqSUEsXR14qpg>;
7   rdf:predicate beol:hasRecipient;
8   rdf:subject <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA> .
9
10  <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/im2jaDG_RKCx999ozN3RZg> a knora-base:TextValue;
11  knora-base:valueHasStandoff
12    <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/im2jaDG_RKCx999ozN3RZg/standoff/3>;
13  knora-base:valueHasString ""
14    Monsieur le Marquis
15    Ayant communiqué à Monsieur Euler [...]"" .
16
17  <http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/im2jaDG_RKCx999ozN3RZg/standoff/3>
18  a beol:StandoffEntityType;
19  beol:standoffEntityType "person";
20  knora-base:standoffTagHasEnd 90;
21  knora-base:standoffTagHasLink <http://rdfh.ch/0801/NbmhfOB1QoGbX3HwXuC3Tg>;
22  knora-base:standoffTagHasStart 90 .
23
24  <http://rdfh.ch/0801/7z7sHo5aTLqSUEsXR14qpg> a beol:person;
25  rdfs:label "Le marquis de Condorcet" .
26
27  <http://rdfh.ch/0801/1PkbSo2nRIeOK9ISp6JSdg> a beol:person;
28  rdfs:label "Anders Johan Lexell" .
29
30  <http://rdfh.ch/0801/NbmhfOB1QoGbX3HwXuC3Tg> a beol:person;
31  beol:hasBirthDate <http://rdfh.ch/0801/NbmhfOB1QoGbX3HwXuC3Tg/values/vtlhBetvSRONLutapTKcaw>;
32  rdfs:label "Leonhard Euler" .
33
34  <http://rdfh.ch/0801/NbmhfOB1QoGbX3HwXuC3Tg/values/vtlhBetvSRONLutapTKcaw> a knora-base:DateValue;
35  knora-base:valueHasCalendar "GREGORIAN";
36  knora-base:valueHasEndJDN 2344633;
37  knora-base:valueHasEndPrecision "DAY";
38  knora-base:valueHasStartJDN 2344633;
39  knora-base:valueHasStartPrecision "DAY";
40  knora-base:valueHasString "1707-04-15 CE" .

```

Listing 2: Query results in the internal schema

```

39  {
40  "id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA",
41  "type": "beol:letter",
42  "beol:creationDate": {
43    "id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/M4eNIJ1MR_eB9mvTGtaZng",
44    "type": "knora-api:DateValue",
45    "knora-api:dateValueHasCalendar": "JULIAN",
46    "knora-api:dateValueHasEndDay": 2,
47    "knora-api:dateValueHasEndEra": "CE",
48    "knora-api:dateValueHasEndMonth": 12,
49    "knora-api:dateValueHasEndYear": 1775,
50    "knora-api:dateValueHasStartDay": 2,
51    "knora-api:dateValueHasStartEra": "CE",
52    "knora-api:dateValueHasStartMonth": 12,
53    "knora-api:dateValueHasStartYear": 1775,
54    "knora-api:valueAsString": "JULIAN:1775-12-02 CE"

```

```

1      },
2      "beol:hasAuthorValue": {
3          "@id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/7ExuBcDbQZyRJRbAcORyKg",
4          "@type": "knora-api:LinkValue",
5          "knora-api:linkValueHasTarget": {
6              "@id": "http://rdfh.ch/0801/1PkbSo2nRIeOK9ISp6JSdg",
7              "@type": "beol:person",
8              "knora-api:arkUrl": {
9                  "@type": "xsd:anyURI",
10                 "@value": "http://ark.dasch.swiss/ark:/72163/1/0801/1PkbSo2nRIeOK9ISp6JSdgr"
11             },
12             "rdfs:label": "Anders Johan Lexell"
13         }
14     },
15     "beol:hasRecipientValue": {
16         "@id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/wCaspGWvS6SZXeuZy3h9MA",
17         "@type": "knora-api:LinkValue",
18         "knora-api:linkValueHasTarget": {
19             "@id": "http://rdfh.ch/0801/7z7sHo5aTLqSueSXR14qpg",
20             "@type": "beol:person",
21             "knora-api:arkUrl": {
22                 "@type": "xsd:anyURI",
23                 "@value": "http://ark.dasch.swiss/ark:/72163/1/0801/7z7sHo5aTLqSueSXR14qpgJ"
24             },
25             "rdfs:label": "Le marquis de Condorcet"
26         }
27     },
28     "beol:hasText": {
29         "@id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/im2jaDG_RKCx999ozN3RZg",
30         "@type": "knora-api:TextValue"
31     },
32     "knora-api:textValueAsXml": "[...] Monsieur le Marquis [...]",
33 },
34 "beol:title": {
35     "@id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/LWf9YstAQ4iq4CI1MmOzUA",
36     "@type": "knora-api:TextValue"
37 },
38 "knora-api:valueAsString": "Annexe 1: Lexell à Condorcet, Saint-Pétersbourg, 2 (13) décembre 1775"
39 },
40 "knora-api:arkUrl": {
41     "@type": "xsd:anyURI",
42     "@value": "http://ark.dasch.swiss/ark:/72163/1/0801/25Ro4c5gSIioLstzJ25FiAC"
43 },
44 "knora-api:hasStandoffLinkToValue": {
45     "@id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA/values/aK3JZSkOR6StGyiaIn6Vlg",
46     "@type": "knora-api:LinkValue",
47     "knora-api:linkValueHasTarget": {
48         "@id": "http://rdfh.ch/0801/NbmhfOB1QoGbX3HwXuC3Tg",
49         "@type": "beol:person",
50         "beol:hasBirthDate": {
51             "@id": "http://rdfh.ch/0801/NbmhfOB1QoGbX3HwXuC3Tg/values/vt1hBetvSRONLutapTKcaw",
52             "@type": "knora-api:DateValue",
53             "knora-api:dateValueHasCalendar": "GREGORIAN",
54             "knora-api:dateValueHasEndDay": 15,
55             "knora-api:dateValueHasEndEra": "CE",
56             "knora-api:dateValueHasEndMonth": 4,
57             "knora-api:dateValueHasEndYear": 1707,
58             "knora-api:dateValueHasStartDay": 15,
59             "knora-api:dateValueHasStartEra": "CE",
60             "knora-api:dateValueHasStartMonth": 4,
61             "knora-api:dateValueHasStartYear": 1707,

```



```

1      "knora-api:valueAsString": "GREGORIAN:1707-04-15 CE"
2    },
3    "knora-api:arkUrl": {
4      "@type": "xsd:anyURI",
5      "@value": "http://ark.dasch.swiss/ark:/72163/1/0801/NbmhfOB1QoGbX3HwXuC3Tg_"
6    },
7    "rdfs:label": "Leonhard Euler"
8  }
9  },
10 "rdfs:label": "Lexell à Condorcet",
11 "@context": {
12   "rdf": "http://www.w3.org/1999/02/22-rdf-syntax-ns#",
13   "knora-api": "http://api.knora.org/ontology/knora-api/v2#",
14   "rdfs": "http://www.w3.org/2000/01/rdf-schema#",
15   "beol": "http://api.dasch.swiss/ontology/0801/beol/v2#",
16   "xsd": "http://www.w3.org/2001/XMLSchema#"
17 }
18 }

```

### Listing 3: Gravsearch response in the complex schema

```

20 {
21   "@id": "http://rdfh.ch/0801/25Ro4c5gSIioLstzJ25FiA",
22   "@type": "beol:letter",
23   "beol:creationDate": {
24     "@type": "knora-api:Date",
25     "@value": "JULIAN:1775-12-02 CE"
26   },
27   "beol:hasAuthor": {
28     "@id": "http://rdfh.ch/0801/1PkbSo2nRIeOK9ISp6JSdg"
29   },
30   "beol:hasRecipient": {
31     "@id": "http://rdfh.ch/0801/7z7sHo5aTLqSUEsXR14qpg"
32   },
33   "beol:hasText": "\n Monsieur le Marquis\n Ayant communiqué à Monsieur Euler [...]",
34   "beol:title": "Annexe 1: Lexell à Condorcet, Saint-Pétersbourg, 2 (13) décembre 1775",
35   "knora-api:arkUrl": {
36     "@type": "xsd:anyURI",
37     "@value": "http://ark.dasch.swiss/ark:/72163/1/0801/25Ro4c5gSIioLstzJ25FiAC"
38   },
39   "knora-api:hasStandoffLinkTo": {
40     "@id": "http://rdfh.ch/0801/NbmhfOB1QoGbX3HwXuC3Tg"
41   },
42   "rdfs:label": "Lexell à Condorcet",
43   "@context": {
44     "rdf": "http://www.w3.org/1999/02/22-rdf-syntax-ns#",
45     "knora-api": "http://api.knora.org/ontology/knora-api/simple/v2#",
46     "rdfs": "http://www.w3.org/2000/01/rdf-schema#",
47     "beol": "http://api.dasch.swiss/ontology/0801/beol/simple/v2#",
48     "xsd": "http://www.w3.org/2001/XMLSchema#"
49   }
50 }

```

### Listing 4: Gravsearch response in the simple schema

```

50 PREFIX knora-api: <http://api.knora.org/ontology/knora-api/simple/v2#>
51 PREFIX foaf: <http://xmlns.com/foaf/0.1/>

```

```

1
2 CONSTRUCT {
3   ?person knora-api:isMainResource true .
4   ?person foaf:familyName ?familyName .
5   ?person foaf:givenName ?givenName .
6 } WHERE {
7   ?person a knora-api:Resource .
8   ?person a foaf:Person .
9   ?person foaf:familyName ?familyName .
10  ?familyName a xsd:string .
11  ?person foaf:givenName ?givenName .
12  ?givenName a xsd:string .
13  FILTER(?familyName = "Euler")
14 }

```

Listing 5: Gravsearch query using FOAF

## Appendix B. A simple Gravsearch query

```

15
16
17
18 PREFIX example: <http://example.org/ontology/0001/example/simple/v2#>
19 PREFIX knora-api: <http://api.knora.org/ontology/knora-api/simple/v2#>
20
21 CONSTRUCT {
22   ?book knora-api:isMainResource true .
23   ?book ?linkProp ?person .
24   ?person example:hasFamilyName ?familyName .
25 } WHERE {
26   ?book a example:Book ;
27   example:hasPublisher <http://rdfh.ch/0001/B3lQa6tSymIq7> ;
28   ?linkProp ?person .
29   ?person example:hasFamilyName ?familyName .
30 }
31 OFFSET 0

```

Listing 6: A Simple Query

## Appendix C. A Gravsearch template

```

32
33
34
35
36 PREFIX beol: <http://beol.dasch.swiss/ontology/0801/beol/simple/v2#>
37 PREFIX knora-api: <http://api.knora.org/ontology/knora-api/simple/v2#>
38
39 CONSTRUCT {
40   ?letter knora-api:isMainResource true .
41   ?letter beol:creationDate ?date .
42   ?letter beol:hasAuthor ?author .
43   ?letter beol:hasRecipient ?recipient .
44 } WHERE {
45   ?letter a beol:letter .
46   ?letter beol:creationDate ?date .
47
48   ?letter beol:hasAuthor ?author .
49   ?author beol:hasIAFIdentifier ?authorIAF .
50   FILTER(?authorIAF = "${iaf1}" || ?authorIAF = "${iaf2}")
51
52   ?letter beol:hasRecipient ?recipient .
53   ?recipient beol:hasIAFIdentifier ?recipientIAF .
54   FILTER(?recipientIAF = "${iaf1}" || ?recipientIAF = "${iaf2}")

```

```

1  }
2  ORDER BY ?date
3  OFFSET ${offset}

```

### Listing 7: A Gravsearch template

#### Appendix D. A user-generated query

```

15
16
17
18
19
20
21
22
23
24
25 PREFIX beol: <http://beol.dasch.swiss/ontology/0801/beol/simple/v2#>
26 PREFIX knora-api: <http://api.knora.org/ontology/knora-api/simple/v2#>
27
28 CONSTRUCT {
29   ?letter knora-api:isMainResource true .
30   ?letter beol:creationDate ?date .
31 } WHERE {
32   ?letter a beol:letter .
33
34   ?letter beol:creationDate ?date .
35   FILTER(?date >= "GREGORIAN:1700-1-1"^^knora-api:Date)
36
37   ?letter beol:hasAuthor <http://rdfh.ch/biblio/Johann_I_Bernoulli> .
38   ?letter beol:mentionsPerson <http://rdfh.ch/biblio/Leonhard_Euler> .
39
40   FILTER NOT EXISTS {
41     ?letter beol:mentionsPerson <http://rdfh.ch/biblio/Daniel_I_Bernoulli> .
42   }
43
44   ?letter beol:hasText ?text .
45   FILTER knora-api:matchText(?text, "Geometria")
46 }
47 ORDER BY ?date
48 OFFSET ${offset}
49
50
51

```

### Listing 8: A user-generated query

## Appendix E. GUI widget

Figure 2. Advanced Search Widget

Extended search ×

Ontology  
The BEOL ontology ▾

Resource Type  
Letter ▾

Property	Comparison Operator	Date
Date of creation ▾	<input checked="" type="checkbox"/> sort criterion since	1 Jan 1700 CE (Gregorian) 📅

Property	Comparison Operator	Resource
Author ▾	is	Johann I Bernoulli ▾
Mentioned person ▾	is	Leonhard Euler ▾
Mentioned person ▾	is not	Daniel (I) Bernoulli ▾

Property	Comparison Operator	Words
Text ▾	<input type="checkbox"/> sort criterion contains the word(s)	Geometria ▾

+   -

×   >

## Appendix F. Searching for text markup

```

PREFIX knora-api: <http://api.knora.org/ontology/knora-api/v2#>
PREFIX standoff: <http://api.knora.org/ontology/standoff/v2#>
PREFIX beol: <http://beol.dasch.swiss/ontology/0801/beol/v2#>

CONSTRUCT {
  ?letter knora-api:isMainResource true .
  ?letter beol:hasText ?text .
} WHERE {
  ?letter a beol:letter .
  ?letter beol:hasText ?text .
  ?text knora-api:textValueHasStandoff ?standoffParagraphTag .
  ?standoffParagraphTag a standoff:StandoffParagraphTag .
  FILTER knora-api:matchTextInStandoff(?text, ?standoffParagraphTag, "Richtigkeit")
}

```

Listing 9: Search for Text Markup

## References

- [1] Bibliothèque nationale de France SPARQL endpoint. <https://data.bnf.fr/current/sparql.html>.

- [2] ISIDORE platform SPARQL endpoint. <https://isidore.science/sparql>.
- [3] Europeana SPARQL endpoint. <http://sparql.europeana.eu/>.
- [4] C. Allocca, A. Adamou, M. d’Aquin and E. Motta, SPARQL Query Recommendations by Example, in: *13th European Semantic Web Conference (ESWC) 2016*, Crete, Greece, 2016. <http://oro.open.ac.uk/46757/>.
- [5] O. Bruneau, E. Gaillard, N. Lasolle, J. Lieber, E. Nauer and J. Reynaud, A SPARQL Query Transformation Rule Language: Application to Retrieval and Adaptation in Case-Based Reasoning, in: *ICCBR 2017: Case-Based Reasoning Research and Development, 25th International Conference on Case-Based Reasoning*, Trondheim, Norway, 2017. <https://hal.inria.fr/hal-01661651>.
- [6] L. Rosenthaler, P. Fornaro and C. Clivaz, DASCH: Data and Service Center for the Humanities, *Digital Scholarship in the Humanities* **30** (2015), i43–i49, Issue suppl\_1. doi:10.1093/lc/fqv051.
- [7] S. Kirrane, A. Mileo and S. Decker, Access Control and the Resource Description Framework: A Survey, *Semantic Web* **8**(2) (2017), 311–352. doi:10.3233/SW-160236.
- [8] O. Hartig and B. Thompson, Foundations of an Alternative Approach to Reification in RDF, *Computing Research Repository (CoRR)* (2014). <https://arxiv.org/abs/1406.3399>.
- [9] T. Schweizer, S. Alassi, M. Mattmüller, L. Rosenthaler and H. Harbrecht, An Interactive, Multi-layer Edition Of Jacob Bernoulli’s Scientific Notebook Meditationes As Part Of Bernoulli-Euler Online, in: *Digital Humanities 2019 Conference Abstracts*, Utrecht, Netherlands, 2019. <https://dev.clariah.nl/files/dh2019/boa/0115.html>.
- [10] D.D. McCarthy, The Julian and Modified Julian Dates, *Journal for the History of Astronomy* **29**(4) (1998), 327–330. doi:10.1177/002182869802900402.
- [11] A. Grafton, Some Uses of Eclipses in Early Modern Chronology, *Journal of the History of Ideas* **64**(2) (2003), 213–229. doi:10.1353/jhi.2003.0024.
- [12] GODOT, Graph of Dated Objects and Texts. <https://godot.date/>.
- [13] E.A. Fellmann, Leonhard Euler, *Historisches Lexikon der Schweiz (HLS)*. <https://hls-dhs-dss.ch/de/articles/018751/2005-11-22/>.
- [14] chronOntology, iDAI.chronontology. <https://chronontology.dainst.org>.
- [15] PeriodO, A gazetteer of periods for linking and visualizing data. <https://perio.do>.
- [16] DBpedia SPARQL endpoint. <http://dbpedia.org/sparql>.
- [17] Wikidata SPARQL endpoint. <https://query.wikidata.org>.
- [18] D. Rogers, The Enduring Myth of the SPARQL Endpoint. <https://daverog.wordpress.com/2013/06/04/the-enduring-myth-of-the-sparql-endpoint/>.
- [19] GraphQL, GraphQL. <https://www.howtographql.com/>.
- [20] HyperGraphQL, HyperGraphQL. <https://www.hypergraphql.org/>.
- [21] S. Klarman, Querying DBpedia with GraphQL. <https://medium.com/@sklarman/querying-linked-data-with-graphql-959e28aa8013>.
- [22] R. Taelman, M.V. Sande and R. Verborgh, GraphQL-LD: Linked Data Querying with GraphQL. <https://comunica.github.io/Article-ISWC2018-Demo-GraphQLD/>.
- [23] T. Ioannidis, G. Garbis, K. Kyzirakos, K. Bereta and M. Koubarakis, Evaluating Geospatial RDF stores Using the Benchmark Geographica 2, 2019. <https://arxiv.org/abs/1906.01933>.
- [24] D. Calvanese, P. Liuzzo, A. Mosca, J. Remesal, M. Rezk and G. Rull, Ontology-based data integration in EPNet: Production and distribution of food during the Roman Empire, *Engineering Applications of Artificial Intelligence* **51** (2016), 212–229. doi:10.1016/j.engappai.2016.01.005.
- [25] D. Calvanese, B. Cogrel, S. Komla-Ebri, R. Kontchakov, D. Lanti, M. Rezk, M. Rodriguez-Muro and G. Xiao, Ontop: Answering SPARQL queries over relational databases, *Semantic Web* **8**(3) (2017), 471–487. doi:10.3233/SW-160217.
- [26] W3C, SHACL. <https://www.w3.org/TR/shacl/>.
- [27] P. Seifer, M. Leinberger, R. Lämmel and S. Staab, Semantic Query Integration With Reason, *The Art, Science, and Engineering of Programming* **3**(3) (2019). doi:10.22152/programming-journal.org/2019/3/13.
- [28] TEI-Consortium (ed.), Guidelines for Electronic Text Encoding and Interchange, Last updated on 13th February 2020. <http://www.tei-c.org/P5/>.