

Knowledge Graphs

Lecture 13: Summary

Markus Krötzsch

Knowledge-Based Systems

TU Dresden, 27 Jan 2026

More recent versions of this slide deck might be available.
For the most current version of this course, see
https://iccl.inf.tu-dresden.de/web/Knowledge_Graphs/en

Overview: KG technologies in this course

We have covered a range of KG-related technologies:

- **SPARQL:** Main query language for knowledge graphs
- **Datalog:** Rule-based query language, can be used with KGs and extended with some SPARQL features (filters, aggregates)
- **OWL:** Ontology language to create terminological models on top of KGs
- **SHACL & ShEx:** Constraint languages to detect problems in KGs

↪ different approaches and goals, but some overlaps in functionality

What SPARQL can do

SPARQL can capture some functionality of the other paradigms:

- **Datalog:** non-recursive Datalog queries and those with simple forms of recursion can be expressed in SPARQL
- **SHACL/ShEx:** SPARQL can be used to describe test queries and competency questions directly (should capture most of SHACL, except undefined recursive features)
- **OWL:** SPARQL is powerful enough to query for entailments from OWL QL ontologies, but larger OWL fragments (even RL and EL) require more powerful recursion that is not available in SPARQL

What Datalog can do

Datalog can capture some functionality of the other paradigms:

- **SPARQL:** when extended with filters and aggregates, many SPARQL queries can be written in Datalog (main exception: OPTIONAL)
- **SHACL/ShEx:** Datalog can be used to describe test queries and competency questions directly (should capture most of SHACL, except recursive features under some semantic interpretations)
- **OWL:** Datalog is powerful enough to compute entailments from OWL QL, OWL RL, and OWL EL ontologies, more complex OWL fragments can be supported using Datalog extensions but performance does not match that of dedicated OWL reasoners

What OWL can do

OWL can capture some functionality of the other paradigms:

- **SPARQL:** OWL can only express tree-shaped queries without advanced features like aggregation
- **Datalog:** OWL RL is the “Datalog fragment” of OWL, but rather limited as rule language
- **SHACL/ShEx:** OWL uses an open-world semantics and is not suitable for constraint validation (OWL axioms entail new information rather than reporting a constraint violations in case of missing information)

What SHACL and ShEx can do

SHACL and ShEX are specific to constraint validation and not suitable for query answering or ontological reasoning.

- Might be least justified as separate standards with own syntax, especially since all scalable implementations are based on translations to SPARQL
- Some features, such as closed shapes, are easier to express than in other formalisms

Knowledge Graphs – What worked, what didn't

Original goal or vision

Distributed global data space (semantic web)

Data exchange format for Web pages

Data exchange format for Web sites/APIs

Data exchange format for databases

Semantic interpretation over syntactic form

Substrate for intelligent (online) agents

Flexible, schema-less data storage and query

Expressive graph queries and analytics

Future-safe graph data format

Was it achieved?

No (never realised to any relevant extent)

Partly (RDFa on many HTML pages, crawled by major engines)

Mostly no (linked open data APIs are a niche, mostly for SPARQL store connectivity)

Partly (CSV and JSON more common)

No (RDF semantics irrelevant in practice)

Mostly no (but agents may also use RDF)

Yes (main goal of KGs in industry)

Partly (specialised tools for advanced analytics, but KGs often more suitable than RDBs)

Yes (strong standards, mature libraries and tools, today's best bet for interoperable graphs)

Consultation

Summary

SPARQL, Datalog, SHACL/ShEx, and OWL have some overlaps in functionality ...

... but each have specific features not supported by the others

What's next?

- No lecture next week
- Examinations