

Alex Ivliev   Markus Krötzsch   Maximilian Marx  
Knowledge-Based Systems Group · TU Dresden

# **SPARQLing Datalog for Rule-Based Reasoning over Large Knowledge Graphs**

**ESWC 2026 · Research Track**

# Combining SPARQL and Datalog

Datalog Engines such as **Nemo** support SPARQL imports

```
labelProp(rdfs:label) .                                % initial fact

@import subPropOf :- sparql {
  endpoint = <https://query.wikidata.org/sparql>,
  query = "SELECT ?x ?y WHERE { ?x rdfs:subPropertyOf ?y }"
} .

labelProp(?x) :- subPropOf(?x, ?y), labelProp(?y) .    % recursive rule
```

# Combining SPARQL and Datalog

Datalog Engines such as **Nemo** support SPARQL imports

```
labelProp(rdfs:label) . % initial fact
```

```
@import subPropOf :- sparql {  
  endpoint = <https://query.wikidata.org/sparql>,  
  query = "SELECT ?x ?y WHERE { ?x rdfs:subPropertyOf ?y }"  
} .
```

```
labelProp(?x) :- subPropOf(?x, ?y), labelProp(?y) . % recursive rule
```

## Problems:

- Its expensive to pre-fetch all the data
- Its difficult to handle *recursive dependencies* between imports and rules

# Combining SPARQL and Datalog

Datalog Engines such as **Nemo** support SPARQL imports

```
labelProp(rdfs:label) . % initial fact
```

```
@import triple :- sparql {  
  endpoint = <https://query.wikidata.org/sparql>,  
  query = "SELECT ?s ?p ?o WHERE { ?s ?p ?o }"  
} .
```

```
labelProp(?x) :- triple(?x, rdfs:subPropertyOf, ?y), labelProp(?y) .
```

## Goals:

- Method to dynamically fetch only the data that is needed
- Optimizations that make that feasible in practice

# Incremental SPARQL Imports

```
@import triple :- sparql { query = "SELECT ?s ?p ?o WHERE { ?s ?p ?o }" } .
```

```
result(?s, ?o) :- triple(?s, ?p, ?o), a(?s, ?z), b(?z, ?o) .
```

import atom                      binding atoms

# Incremental SPARQL Imports

```
@import triple :- sparql { query = "SELECT ?s ?p ?o WHERE { ?s ?p ?o }" } .
```

```
result(?s, ?o) :- triple(?s, ?p, ?o), a(?s, ?z), b(?z, ?o) .
```

import atom                      binding atoms

**Bindings** ( $\pi_{?s, ?o}(a \bowtie b)$ ):

?s	?o
alice	london
bob	london
bob	berlin

# Incremental SPARQL Imports

```
@import triple :- sparql { query = "SELECT ?s ?p ?o WHERE { ?s ?p ?o }" } .
```

```
result(?s, ?o) :- triple(?s, ?p, ?o), a(?s, ?z), b(?z, ?o) .
```

import atom                      binding atoms

**Bindings** ( $\pi_{?s, ?o}(a \bowtie b)$ ):

?s	?o
alice	london
bob	london
bob	berlin

**SPARQL query:**

```
SELECT ?s ?p ?o WHERE {  
  ?s ?p ?o .  
  VALUES (?s ?o) {  
    (alice london)  
    (bob london)  
    (bob berlin)  
  }  
}
```

# Incremental SPARQL Imports

```
@import triple :- sparql { query = "SELECT ?s ?p ?o WHERE { ?s ?p ?o }" } .
```

```
result(?s, ?o) :- triple(?s, ?p, ?o), a(?s, ?z), b(?z, ?o) .
```

import atom                      binding atoms

**Bindings** ( $\pi_{?s, ?o}(a \bowtie b)$ ):

?s	?o
alice	london
bob	london
bob	berlin
carol	paris

**SPARQL query:**

```
SELECT ?s ?p ?o WHERE {  
  ?s ?p ?o .  
  VALUES (?s ?o) {  
    (carol paris)  
  }  
}
```





# Cartesian Products

```
@import triple :- sparql { query = "SELECT ?s ?p ?o WHERE { ?s ?p ?o }" } .
```

```
result(?s, ?o) :- triple(?s, ?p, ?o), a(?s, ?x), b(?x, ?y), c(?o, ?z) .  
import atom group 1 group 2
```

## Bindings 1

<u>?s</u>
alice
bob

## Bindings 2

<u>?o</u>
london
berlin

## SPARQL query:

```
SELECT ?s ?p ?o WHERE {  
  ?s ?p ?o .  
  VALUES (?s) { (alice) (bob) }  
  VALUES (?o) { (london) (berlin) }  
}
```

# Cartesian Products

```
@import triple :- sparql { query = "SELECT ?s ?p ?o WHERE { ?s ?p ?o }" } .
```

```
result(?s, ?o) :- triple(?s, ?p, ?o), a(?s, ?x), b(?x, ?y), c(?o, ?z) .  
import atom group 1 group 2
```

## Bindings 1

<u>?s</u>
alice
bob
<b>carol</b>

## Bindings 2

<u>?o</u>
london
berlin
<b>paris</b>

## Variant 1:

```
SELECT ?s ?p ?o WHERE {  
  ?s ?p ?o .  
  VALUES (?s) { (carol) }  
  VALUES (?o) {  
    (london)  
    (berlin)  
    (paris)  
  }  
}}
```

## Variant 2:

```
SELECT ?s ?p ?o WHERE {  
  ?s ?p ?o .  
  VALUES (?s) {  
    (alice) (bob)  
  } VALUES (?o) {  
    (paris)  
  }  
}
```

# Example Applications

Handcrafted examples combining Datalog and SPARQL imports



## Common Ancestors

Find the most recent common ancestors of two people on Wikidata




## Winning Streaks

Compute the longest winning streaks in a series or league of sports events



## Constraint Violations

Detect disjoint-class, inverse-property, and none-of violations in Wikidata



## Open Access

Find Oxford employees who published in open-access journals

# Evaluation

## System Comparison

Runtime on all benchmarks

**Nemo v0.10.0:** new implementation

**Nemo v0.8.0:** no optimizations

**VLog:** incremental imports


## Ablation Study

Disable individual optimizations

- Merging import atoms
- Avoiding Cartesian products

Timeout: 10 min • 16 GiB memory limit • Median of 10 runs

# Evaluation: System Comparison



## Evaluation: Ablation Study

Test impact of different optimizations

Benchmark	Full	No Merging	No Cart. Prod.
Common Ancestors	24 s	29 s	32 s
Winning Streaks	1 s	42 s	1 s
Disjoint Classes	32 s	timeout	25 s
Inverse Constraint	48 s	timeout	44 s
Non-of Constraint	10 s	timeout	9 s
Open Access	14 s	67 s	timeout

Every optimization contributes to performance gains

# Summary

- New method for on-demand SPARQL imports in rule languages
- Optimizations that improve performance significantly in practice
- Try Nemo at: `tools.iccl.inf.tu-dresden.de/nemo`



The screenshot shows the Nemo Graph Rule Engine interface. At the top left is the Nemo logo (a stylized 'N' in a circle) and the text 'Nemo Graph Rule Engine'. To the right are links for 'Examples' and 'Help'. The main area contains a SPARQL query editor with the following code:

```
1  %! Find common ancestors of Ada and Moby.
18 @parameter $personId2 = wd:Q14045 . % Moby
19
20 % Import predicate "wdParent" (father or mother) through SPARQL:
21 @import wdParent :- sparql{
22   endpoint = <https://query.wikidata.org/sparql>,
23   query = """
24     PREFIX wdt: <http://www.wikidata.org/prop/direct/>
25     SELECT ?child ?parent WHERE { ?child (wdt:P22|wdt:P25) ?parent }
26     """
27 } .
28 % Import predicate "wdLabel" (English label) through SPARQL:
29 @import wdLabel :- sparql{
30   endpoint = <https://query.wikidata.org/sparql>,
31   query = """
32     PREFIX wikibase: <http://wikiba.se/ontology#>
33     SELECT ?qid ?qidLabel WHERE {
34       SERVICE wikibase:label {
```

At the bottom left, there is a blue play button icon and a text input field containing 'Add input files'. On the right side of the interface, there is a QR code.