

Talk2PowerSystem

Democratizing Power System Analytics via Generative AI

Vladimir Alexiev
Industrial semantic interoperability summit 2026
DNV, Oslo, 9-11 June 2026



Statnett is the transmission system operator in the Norwegian energy system. Our mission is to ensure a reliable power supply by building and operating the transmission grid. Statnett is a state enterprise owned by the Norwegian state through the Ministry of Energy


Graphwise enables organizations to unlock return on investment for enterprise AI by delivering the most comprehensive and trusted industry solution in the field of knowledge graphs and semantic AI technologies

Unicus/auticon is an IT consultancy, all of our consultants have a neurodiverse background. We were established in Norway in 2009 and are now represented in 15 countries across the globe. Our services also include inclusion services, leveraging our experience in working with autistic talents

"The most interesting project in my life"
-- Andreas Kimsas, project owner, Statnett



Graphwise in a Nutshell



 info@graphwise.ai

 www.graphwise.ai

 [linkedin.com/company/graphwise/](https://www.linkedin.com/company/graphwise/)


*Read more: <https://www.gartner.com/peer-insights/home> (4.6 is the average of the GraphDB and PoolParty ratings)

One Stop Shop for Graph AI

Graphwise enables enterprises to

- connect the dots across data sources
- surface knowledge
- propel decision making

We leverage graphs to combine **domain knowledge** and diverse data into a **holistic view** that enhances AI's ability to derive meaningful **trustworthy insights**



T2PS Project

- Start Feb 2025, end Apr 2026 (14m)
- Monthly sprints, quarterly milestones

Work package / Task	Start-end	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15
WP1 Semantic Infrastructure and Integration																
T1.1 Semantic Infrastructure	M1-M14		1.1												1.1	
T1.2 CIM Semantic Data	M1-M5					1.2										
T1.3 Semantic Integration of Other Data	M3-M11											1.3				
WP2 LLM, NLQ, Analytics																
T2.1 KGQA Dataset, Evaluation and Validation	M1-M11					2.1										
T2.2 LLM NLQ, RAG, Analytics	M3-M11					2.2						2.2				
T2.3 ML Operations	M3-M11											2.3				
WP3 BIM-CIM Integration																
T3.1 BIM-CIM Research, Requirements, Design	M3-M8					6.3			6.4							
T3.2 BIM-CIM Conversion and Integration	M6-M11											3.1				
WP4 Software Development																
T4.1 Develop Consumption APIs and Chatbot	M9-M14											4.1			4.1	
T4.2 Develop CIM Model Visualizations	M12-M14														4.2	
T4.3 Integration in Enterprise IT Infrastructure	M12-M15															4.3
WP5 Research Collaboration and Dissemination																
T5.1 Ongoing LLM Research	M1-M11															
T5.2 Data Spaces and Research Proposals	M9-M15											5.3				5.4
T5.3 Dissemination and Communication	M1-M15		5.1									5.2	5.5			5.1
T5.4 Knowledge Transfer	M15-M15															
WP6 Project Management, Requirements, Agile...																
T6.1 Project Management	M1-M15		6.1			6.2			6.2			6.2				6.2
T6.2 Requirements, Design, Client Relations	M1-M14		6.3			6.3			6.4							
		Milestone 1		Milestone 2		Milestone 3		Milestone 4		Milestone 5						
		Discovery		First version		Second version		Third version		Final delivery						

Open results: [source](#), [Q&A dataset](#), [dissemination](#)

- [Graphwise: Semantic Digital Twins Use Case](#)
- [CIMon Chatbot](#): ask your own questions. See technical details in the "Components" page
- [Talk2PowerSystem Promo Video](#): explanation/motivation video
- [CIMon Demo Video](#): actual interaction with the bot (narration in final webinar)
- [CIM GraphDB repository](#) (semantic database) with public data (Nordic44 for TSO and Telemark120 for DSO)
- [Talk2PowerSystem Project wiki](#) including:
 - [Repositories related to the project](#)
 - [KGQA Corpus Generation](#)
 - [Evaluation Results](#)

Main Components

- Electrical CIM/CGMES/NCP, subsetting inference
- Public electric datasets: Nordic44 (TSO), Telemark120 (DSO), some timeseries (Cognite)
- Q&A Dataset generation
- Chatbot capabilities and tools: IRI discovery, n-shot learning, SPARQL querying, timeseries querying
- Diagrams: electrical, vizGraph, soon geospatial; interactivity
- Evaluation framework: automated scoring, error analysis
- Local LLM fine-tuning

Why LLMs need KGs

GenAI brings big opportunities, but also big challenges

1. **LLMs** let enterprises exploit data and knowledge in a better way

- Business users can talk to data directly and get instant insights
- Extracting relevant information from documents and digital assets got **10x easier**
- Building and using knowledge graphs becomes much easier


2. **LLMs** need grounding context to be successful and useful

- Context = Diverse factual data + Rich domain knowledge
- Knowledge graphs provide data with structure and unify meaning across diverse sources, enabling LLMs to understand relationships and accurately answer complex questions
- Feeding LLMs with proprietary data is the only way to get competitive insights

3. Reliability and quality are crucial to bringing **GenAI** to production

- **Limit hallucinations** by using graphs to educate LLMs and assess accuracy for specific tasks and topics
- A comprehensive enterprise knowledge graph platform reduces the integration cost and burden

Chatbot Architecture Overview




Reasoning: a Difficult CIM Query

- List all substations connected by a Line to "ARENDAL"

```
PREFIX cim: <https://cim.ucaiug.io/ns#>
PREFIX sesame: <http://www.openrdf.org/schema/sesame#>
select ?sub1Name ?lineName ?sub2Name {
  {select distinct * {
    values ?sub1Name {"ARENDAL"}
    ?sub1 a cim:Substation;
    cim:IdentifiedObject.name ?sub1Name;
    (cim:EquipmentContainer.Equipments|cim:Substation.VoltageLevels|cim:VoltageLevel.Bays)+ / # equipment in ?sub1
    cim:ConductingEquipment.Terminals / cim:Terminal.ConnectivityNode /
    cim:ConnectivityNode.Terminals / cim:Terminal.ConductingEquipment / # connected to segment
    cim:Equipment.EquipmentContainer ?line. # part of ?line
    ?line a cim:Line; cim:IdentifiedObject.name ?lineName}}
  {select distinct * {
    ?sub2 a cim:Substation;
    cim:IdentifiedObject.name ?sub2Name;
    (cim:EquipmentContainer.Equipments|cim:Substation.VoltageLevels|cim:VoltageLevel.Bays)+ / # equipment in ?sub2
    cim:ConductingEquipment.Terminals / cim:Terminal.ConnectivityNode /
    cim:ConnectivityNode.Terminals / cim:Terminal.ConductingEquipment / # connected to segment
    cim:Equipment.EquipmentContainer ?line}}
  filter(?sub1 != ?sub2)
}
```

The query is complex, hard for LLM to generate:

- 3 kinds of Container-Equipment relations
- Some containers are optional (eg Bay in Substation)
- Line is also considered a container (of ACLineSegments)
- From Substation to Line: Container-Equipment-Terminal-ConnectivityNode-Terminal-Equipment-Container
- Then up the other side




Add Reasoning (Inference)

Inferred relations, namespace **cimr**: (CIM Rules)

- **Inverse**: in CIM, all relations have inverses.
- **Union**:
`cim:EquipmentContainer.Equipments|cim:Substation.VoltageLevels|cim:VoltageLevel.Bays -> cimr:hasPart; inverse cimr:isPart`
- **Union**:
`cim:Terminal.ConductingEquipment|cim:Terminal.AuxiliaryEquipment -> cimr:Terminal.Equipment; inverse cimr:Equipment.Terminals`
- **Transitive closure**: `cimr:hasPart+ -> cimr:hasPartTransitive; inverse cimr:isPartTransitive`
- **Property path**: `cim:ConductingEquipment.Terminals / cim:Terminal.ConnectivityNode / cim:ConnectivityNode.Terminals / cim:Terminal.ConductingEquipment -> cimr:connectedTo (symmetric)`
- **Property path**: `cimr:hasPartTransitive / cimr:connectedTo / cimr:isPartTransitive -> cimr:connectedThroughPart (symmetric)`


Dependency Graph of Relations



Simplified Query

After adding `cimr:` derived relations, the query becomes much simpler:

```
select ?sub1Name ?lineName ?sub2Name {  
  values ?sub1Name {"ARENDAL"}  
  ?sub1 a cim:Substation; cim:IdentifiedObject.name ?sub1Name;  
    cimr:connectedThroughPart ?line.  
  ?line a cim:Line; cim:IdentifiedObject.name ?lineName.  
  ?sub2 a cim:Substation; cim:IdentifiedObject.name ?sub2Name;  
    cimr:connectedThroughPart ?line.  
  filter(?sub1 != ?sub2)  
}
```



Can you give me an example of a Generating unit that is connected to a substation?



Yes—here's one example from the dataset:

- **GeneratingUnit:** BLAFALLIG2
IRI: <urn:uuid:f1769964-9aeb-11e5-91da-b8763fd99c5f>
- **Substation:** BLAFALLI
IRI: <urn:uuid:f176968e-9aeb-11e5-91da-b8763fd99c5f>

This is based on the generating unit being contained in the substation via the transitive "has part" relationship (`cimr:hasPartTransitive`).

SPARQL used:

```
SELECT ?gu ?guName ?sub ?subName
WHERE {
  ?sub a cim:Substation ;
       cimr:hasPartTransitive ?gu .
  ?gu a cim:GeneratingUnit .
  OPTIONAL { ?gu cim:IdentifiedObject.name ?guName }
  OPTIONAL { ?sub cim:IdentifiedObject.name ?subName }
}
```


This is a demo version and we cannot guarantee the accuracy of results, nor the availability or security of the system. Please don't include any kind of sensitive data in your questions

Could you show me a Network Area Diagram of the substations and lines that are contained in the model?



Here is a **Network Area Diagram (NAD)** showing the **substations/voltage levels and the lines/transformers between them** for the whole model:


- **Diagram:** PowSyB1-NAD-all.svg
- **IRI:** urn:uuid:4585f674-8011-4dc1-999c-d374b5174c5f

If you want, I can also show a NAD **for a specific area** (e.g., SE1/SE2/SE3/SE4, NO1/NO3/NO4/NO5, FI1) instead of the full model.

 Generated diagram



Ask anything





 Chatbot

 Components

 Predefined questions

- connected to the node ABC.
- Terminals connected to equivalent injection ABC.
- Topological islands sharing angle reference with node ABC.
- Synchronous areas in the system.
- Connectivity nodes within voltage level XYZ.
- Topological node interacting with voltage level XYZ.

 Operational & Control

This is a demo version and we cannot guarantee the accuracy of results, nor the availability or security of the system. Please don't include any kind of sensitive data in your questions

HELSINKI (Substation):

- **PowSyBI Single-Line-Diagram of substation HELSINKI**
Link/file: PowSyBI-SLD-substation-HELSINKI.svg
Diagram IRI: urn:uuid:ad4c22dd-e4f4-40e1-a1e4-83a0f2e4b889

If you'd like an even *wider* view that includes how Helsinki ties into neighbors, there are also multi-substation SLDs:


- **HELSINKI + ESTLINK_HVDC:** PowSyBI-SLD-2substations-HELSINKI-and-ESTLINK_HVDC.svg
- **HELSINKI + OULU:** PowSyBI-SLD-2substations-HELSINKI-and-OULU.svg
- **HELSINKI + VYBORG_HVDC:** PowSyBI-SLD-2substations-HELSINKI-and-VYBORG_HVDC.svg

Which one do you want to look at next?

 [Generated diagram](#) 



Ask anything





Chatbot

Components

Predefined questions

Equipment & Components

- Breaker equipment associated with current transformer ABC.
- Load break switches.
- Linear shunt compensators part of voltage level XYZ and transitive parts of substation ABC.
- Power transformer ends, part of terminal ABC with resistance Y.
- Potential transformers connected via terminals to breaker ABC.
- Rotating machine activities connected to grid state alteration ABC.

This is a demo version and we cannot guarantee the accuracy of results, nor the availability or security of the system. Please don't include any kind of sensitive data in your questions

comparison.

Can you list all ACLines in the model and the corresponding nominal voltages of these lines, with voltages specified in one colum for one end and another colum for the other end?

Yes. Below is a 2-end listing of **all cim:ACLineSegment** objects in the model, with the **nominal voltage at each end** shown in separate columns.

Note: In the raw data, some nominal voltages appear with different RDF literal datatypes/lexical forms (e.g., 135 vs 132 vs 4.0E-1), which can cause duplicate-looking rows for the same line in a SPARQL result if you don't normalize/cast. The table below reflects the values as stored.

ACLineSegment (name)	End 1 nominal voltage (kV)	End 2 nominal voltage (kV)
023 ZT1-EFC ACLS 1	0.23	0.23
023 ZT1-EFC ACLS 2	0.23	0.23
04 TELEMA 2 ACLS 3	0.4	0.4
04 TELEMA2 ACLS1	0.4	0.4
04 TELEMA2 ACLS10	0.4	0.4
04 TELEMA2 ACLS11	0.4	0.4
04 TELEMA2 ACLS12	0.4	0.4
04 TELEMA2 ACLS13	0.4	0.4
04 TELEMA2 ACLS14	0.4	0.4
.....	-	-

Ask anything



Chatbot

Components

Predefined questions

Equipment & Components

- Breaker equipment associated with current transformer ABC.
- Load break switches.
- Linear shunt compensators part of voltage level XYZ and transitive parts of substation ABC.
- Power transformer ends, part of terminal ABC with resistance Y.
- Potential transformers connected via terminals to breaker ABC.
- Rotating machine activities connected to grid state alteration ABC.

This is a demo version and we cannot guarantee the accuracy of results, nor the availability or security of the system. Please don't include any kind of sensitive data in your questions

and I'll return a normalized table with exactly one row per ACLineSegment.

Can you make the same query, but only highlight the ACLines that have different voltages (more than 1% difference)



Below are the **ACLineSegments where the nominal voltage differs by more than 1% between the two ends** (computed as $\frac{|V_1-V_2|}{(V_1+V_2)/2} > 0.01$).

Because of how the voltages are stored (some as xsd:integer, some as xsd:float), the raw SPARQL result can produce multiple rows per same line; **the distinct "problem" lines are these 4:**

ACLineSegment	End1 kV	End2 kV	ΔkV	Δ%
132 ARENDAL-FROLAND ACLS 1	135	132	3.0	2.247191%
132 ARENDAL-T1 ACLS 1	135	132	3.0	2.247191%
132 FROLAND-T_ENGENE ACLS 1	135	132	3.0	2.247191%
132 T_ENGENE-ENGENE ACLS 1 1	135	132	3.0	2.247191%

If you want, I can rerun the query in a way that **deduplicates to exactly one row per ACLineSegment** (e.g., by selecting MIN/MAX voltage per line end or by aggregating per line).



Hint: you can also ask the model, for example How did you derive this answer?

Ask anything

📤



Timeseries Querying/Analysis

Hi, I would like to get the maximum power price for the each day in March 2026, for the price area No1 and No2. Listed as a table

Here is the table with the maximum daily power price for March 2026 for NO1 and NO2.

Date	NO1 max price (EUR/MWh)	NO2 max price (EUR/MWh)
2026-03-01	95.87	95.87
2026-03-02	267.95	239.44
2026-03-03	130.89	135.65
2026-03-04	128.66	133.82
2026-03-05	189.64	192.92
2026-03-06	201.50	206.17
2026-03-07	149.58	152.78
2026-03-08	135.49	137.91
2026-03-09	216.60	220.21
2026-03-10	220.43	225.20
2026-03-11	123.23	125.99
2026-03-12	114.03	109.83
2026-03-13	116.94	119.97
2026-03-14	136.80	139.56
2026-03-15	117.14	119.16
2026-03-16	143.72	146.67
2026-03-17	123.99	126.69
2026-03-18	137.36	141.44
2026-03-19	177.75	182.85
2026-03-20	177.68	182.93

Called Retrieve datapoints from Cognite: Retrieve datapoints for one or more time series

Arguments:

```
{
  "external_id": "Kraftpriser_NO2_13",
  "limit": null,
  "start": "2026-03-01T00:00:00Z",
  "end": "2026-04-01T00:00:00Z",
  "aggregates": "max",
  "granularity": "1day"
}
```

The exact same setup can be used on CIMon Confidential to access > 100 000 timeseries on analogue and discrete measurements, e.g. active power and breaker positions

Two more tricks

Fine Tuning:

Fine tuning: Fine-tuning language models refers to taking a specific (open source) model with existing weights as a baseline and performing additional training to change it's behavior.

SOTA on Local LLMs: state of the art for locally deployable LLMs



N-shot Learning:

Different answers to frequently asked questions are put into a database and the agent tries to understand whether the specific prompt has a similar query in the database.


The screenshot shows the GraphDB web interface. The main panel displays the 'Talk2PowerSystem Q&A Dataset Ontology'. It includes a description, a source URL, and a table of query results. The table has columns for 'subject', 'predicate', and 'object'. The results show four rows of data related to the ontology's creation and modification.

	subject	predicate	object
1	qa:	dct:created	"2025-11-14""xsd:date
2	qa:	dct:creator	https://graphwise.ai
3	qa:	dct:license	https://creativecommons.org/licenses/by-sa/4.0
4	qa:	dct:modified	"2025-11-14""xsd:date


Evaluation Results

- [Evaluation iteration #3 \(M5, April 2026: GPT 5.4\)](#)
- [Evaluation iteration #2.2 \(M5, March 2026: GPT 5.2, 5.4\)](#)
- [Evaluation iteration #2.1 \(M4, Jan 2026: GPT 5.2\)](#)
- [Evaluation iteration #2 \(M3, Sep 2025: GPT 4o, 4.1, 5, GPT-oss, Qwen3\)](#)
 - [Results](#)
 - [Summary/Analysis](#)
 - [Error Analysis](#)
 - [Autocomplete Tool Not Used](#)
 - [KGQA Corpus Generation Problem](#)
 - [Evaluation Framework Erroneously Detects an Error](#)
 - [Autocomplete Tool Fails](#)
 - [Data Problems in the Knowledge Graph](#)
 - [Problem with Ontology Understanding/Use](#)
 - [Summary](#)
 - [Open-source LLMs](#)
- [Evaluation iteration #1 \(M2, Jun 2025: GPT 4o, 4.1\)](#)
 - [Results](#)
 - [Error analysis](#)
- [User-identified Bad Conversations](#)
 - [Inability to generalize knowledge about switches](#)
 - [Couldn't count transformers in substation](#)
 - [Obtain rated currents and nominal voltage for a given Substation](#)
 - [Fails to use inferred properties when helpful/required](#)
 - [Nearly perfect conversation](#)
 - [Synchronous machines and generating units](#)

M5, April 2026, GPT 5.4



Evaluation Results for Open Models (Qwen vs GPT)



Final Webinar

- Talk2PowerSystem final project webinar (CIMon results). Live Demo. Expert Panel with PNNL & Siemens Energy. 21 May 2026, online (and available on demand).
- [Talk2PowerSystem Project Introduction, CIM, Knowledge Graphs & LLMs](#). Andreas Kimsås, Svein Harald Olsen, Vladimir Alexiev
- [Talk2PowerSystem Technical Deep Dive](#). Can LLMs answer complex power system questions? Vladimir Alexiev, Neli Hateva, Andrey Tagarev.
- [CIM-based Data Agents for Interoperability](#). Enabling Control Room AI Applications with CIM and Graph Databases. Alexander Anderson, PNNL
- From FAIR to Enterprise Knowledge Graphs over a Data Mesh. Alexander Garcia-Castro, Siemens Energy (presentation and recording are not available yet)
- [Talk2PowerSystem Final Webinar Q&A Session](#)
- [Talk2PowerSystem Final Webinar Recording](#): available on demand with free registration
- [Connecting Minds: AI Use Cases to Bridge Power Systems and Large Language Models for Practical Applications](#) (PNNL-38003, May 2025): Compilation of 20 high-impact use cases of LLMs for planning, operations, and asset management.

What's next

- Try on real production data
- Integrate into production applications
- Deploy on Local LLM for confidentiality and sovereignty
- Add more competency questions and tools
- Establish a collaboration with more electrical enterprises



Establish a partnership

- Our industry face the same challenges
- Lets solve industry problems at industry scale !!



Test on real power system data

- Statnett expert users
- R&D platform



Local LLMs

- Extending the Q&A dataset
- Tuning medium sized LLM



CimOn for explainability

- Its the holy grail of power system analysis
- Is it within reach?

Call for Collaboration

