


Data and Knowledge Management

MSc topics 2026


Universitetet i Oslo

25 September 2026

Data and Knowledge Management



Logic,
Reasoning



Data Science
AI/ML



Digital
Twins



Semantic Web
Technologies



Data
Management

DKM MSc topics

DKM:

<https://www.mn.uio.no/ifi/studier/masteroppgaver/dkm/>

Reasoning about Descriptions

- An identifier: “King Håkon VIII”

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”
- In Description Logic: $\epsilon(\exists \text{kingOf}.\{\text{Norway}\})$

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”
- In Description Logic: $\epsilon(\exists \text{kingOf}.\{\text{Norway}\})$
- Recent MSc thesis on this topic

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”
- In Description Logic: $\epsilon(\exists \text{kingOf}.\{\text{Norway}\})$
- Recent MSc thesis on this topic
- Remains to be done:

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”
- In Description Logic: $\epsilon(\exists \text{kingOf}.\{\text{Norway}\})$
- Recent MSc thesis on this topic
- Remains to be done:
 - Implementation, extending some existing DL theorem prover X

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”
- In Description Logic: $\epsilon(\exists \text{kingOf}.\{\text{Norway}\})$
- Recent MSc thesis on this topic
- Remains to be done:
 - Implementation, extending some existing DL theorem prover X
 - Depending on X some theory to fill the gap

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”
- In Description Logic: $\epsilon(\exists \text{kingOf}.\{\text{Norway}\})$
- Recent MSc thesis on this topic
- Remains to be done:
 - Implementation, extending some existing DL theorem prover X
 - Depending on X some theory to fill the gap
 - ~~Some theory (not implementation-related)~~

Reasoning about Descriptions

- An identifier: “King Håkon VIII”
- A description: “The king of Norway”
- In Description Logic: $\epsilon(\exists \text{kingOf}.\{\text{Norway}\})$
- Recent MSc thesis on this topic
- Remains to be done:
 - Implementation, extending some existing DL theorem prover X
 - Depending on X some theory to fill the gap
 - ~~Some theory (not implementation-related)~~

Contact: Martin Giese (martingi@uio.no)

Supervisor: Leif Harald Karlsen – leifhka@ifi.uio.no

Supervisor: Leif Harald Karlsen – leifhka@ifi.uio.no

- DATASTAR



[illegible]

What can (and cannot) we learn with GNNs?

- Graph Neural Networks (GNNs), including Graph Transformers (GTs):
 - a family of modern NN architectures
 - work with structured data (e.g., knowledge graphs)
- (Theoretical) expressivity:
 - different GNN architectures express different sets of functions on graphs
 - the landscape is still not complete
 - Example question:
Which logic-based functions can GCNs express?
- (Practical) trainability:
 - "can express in theory" does not imply "can be trained"
 - the landscape is much less clear
 - Example question:
Which logic-based functions can GTs can be trained for?

Supervisors:


Egor V. Kostylev (egork@uio.no), Lorenzo Mastrandrea (lorenzma@uio.no)




- xAI for sleep related breathing disorders
 - Infants
 - Severe cases requiring mechanical ventilation
- Potential topics:
 - Empirical evaluation of explanations
 - Testing xAI approaches
 - Knowledge graphs for explanations
- We design the thesis topic together with you!
- Contact: Thomas Plagemann,
plageman@ifi.uio.no

Reliable Actionable Recourse through Explainable AI

Same predictive performance — different recourse. Which recommendation should we trust?



Multiagent RAG Systems for Industrial Applications



- RAG: Retrieval Augmented Generation, LLMs+structured data
- Project: use multi-agent methods to improve RAG
- Context is industrial data in Bosch

Contact: Martin Giese (martingi@uio.no), Evgeny Kharlamov (evgenykh@uio.no)

Build AI agents that understand industrial data

MSc thesis projects on IMF + LLM agents
IM-Lab (Dept. of Informatics, UiO) with Bosch Research
Twelve topics · real industrial models

25 September 2026




Conceptual models: the key enabler of industrial AI

Palantir – the Ontology



Example of the model



Objects, links and actions of the business – AI agents act on that model.

Revenue 2025 ≈ €4.0 B · market value ≈ €405 B


€405 B

Palantir market value, Sep 2026

Cognite (Oslo) – knowledge graph



Example of the model



Assets, sensors, work orders, documents, 3D – one linked graph of the plant.

Bought by Schneider Electric for ≈ €2.7 B (2026)

€2.7 B

Schneider Electric pays for Cognite, Jun 2026


€19 B → €132 B

digital-twin market, 2025 → 2030

SAP – Knowledge Graph for Joule



Example of the model



Business objects and processes, grounded in the S/4HANA data – Joule agents use it.

Revenue 2025 €36.8 B · Joule customers ×9

€0.9 B → €6.1 B

knowledge-graph market, 2024 → 2030

The model is the moat: AI can only act on what is described precisely – IMF is that description for engineering.

Model examples are schematic, in each vendor's own terms (Palantir Ontology docs; Cognite Data Fusion data model; SAP Knowledge Graph). Figures: Palantir FY2025 10-K; companiesmarketcap.com (23 Sep 2026); Bloomberg/SiliconANGLE 30 Jun 2026; SAP FY2025; MarketsandMarkets 2025; 1.136 USD/EUR. Logos are trademarks of the companies.

From scattered records to one model – the gas-cooling case

1 · One pump, five records – all of them right

As planned (intended)



in the diagram



design drawing



data sheet: **30 kW**



installed: **37 kW**



in the plant


30 ≠ 37 kW – both true, different views

An AI that only reads the documents mixes the records: “the pump has 30 kW and 37 kW” – wrong answers, no warning.


Same as you in Studentweb, Canvas and your bank app: several records of one person – nobody merges them blindly.

2 · The model (IMF): every record gets its view – and links

PRODUCT – what is installed




LOCATION – where it is





part of

c



has part

d



blocks with tag IDs, each in an aspect · dashed = same thing in another aspect · FUNCTION aspect not shown

Gas alarm at A-P02-GD001B – which equipment nearby handles gas?

The agent walks a → b → c → d → e (two queries): **A23HJ105**, terminals **N3** and **N4** – checkable step by step, with every source.

IMF labels every record with its view (planned or real · function, product, location) and links them – on that model, agents can use the data safely.

Pictures: A. Waaler, UiO, 15 Sep 2026. Model: imfid.org – IMF gas cooling system case (2025), excerpt redrawn; RDF and SPARQL in the notes.

Twelve thesis topics – four ways to help agents use the model

Build the model 	Ask the model 	Check the model 	Trust the answers 
1 · Documents → IMF models an agent drafts the model from data sheets and drawings	2 · Ask the plant plain questions → queries over the right views	3 · Consistency checker finds 30 kW vs 37 kW – and explains it	6 · Answers you can trust every answer with source and confidence
4 · Modelling copilot an assistant that speaks the engineer's language	8 · Tool-using agent IMF, SPARQL, SHACL as tools (MCP)	9 · Generate, verify, repair LLM proposes, reasoner checks, agent fixes	10 · Know when you don't know conformal prediction – abstain, don't guess
7 · An LLM that speaks IMF RAG vs fine-tuning on the IMF manual	12 · IMF-QA benchmark questions, traps and “unanswerable” cases	5 · Standards bridge IMF ↔ IDO, AAS, ISO/IEC 81346	11 · Explanations that convince do explanations help engineers catch errors?

All topics: real industrial models, LLM agents + knowledge graphs, co-supervised by the IM-Lab and Bosch Research. Own ideas welcome.

Topics are proposals – shape them with us. Details on the IM-Lab student-topics page (QR code on the last slide).

Join us: how it works and how to apply

What you work on



- Real IMF models from industrial cases (open spec: imfid.org)
- LLM agents with tools, knowledge graphs (RDF, OWL, SHACL, SPARQL), Python
- One of the twelve topics – or your own idea on IMF + AI
- Long (60 ECTS) or short (30 ECTS) thesis – both fit

What you get



- **Co-supervision:** IM-Lab (UiO) and Bosch Research
- Industrial problems from energy (IMF projects), manufacturing (Bosch) and our EU projects
- A team to work with – and a result aimed at a workshop or conference paper
- Open-source LLMs and agent frameworks – no black boxes

How to apply



- Talk to us right after this session
- Or e-mail a short CV, your transcript and your top two topics
- arild@ifi.uio.no (Arild Waaler, head of the IM-Lab)
- **IM-Lab open day (date to be announced):** the full presentation of the lab and the topics – you are invited
- **Scan the code:** all topics and details on the IM-Lab student-topics page



Scan for the topics page:
mn.uio.no/ifi/studier/masteroppgaver/psy/im-lab-student-topics.html

The IM-Lab team – you get a team, not one supervisor



A. Waaler



M. Skjæveland



E. Kamburjan



B. Zhou



Y. Qu



D. Kyritsis



F. Psarommatis



M. Doroudi M.



D. Cameron




E. Kharlamov

Pick a topic today – and come to the IM-Lab open day for the full picture (date to be announced).

Photo: SIRIUS presentation files. Team: mn.uio.no/ifi/english/research/networks/im-lab; photos from the UiO people pages. QR: IM-Lab student-topics page on mn.uio.no.

Which uncertain matches should an LLM examine first?




Ontologies describe concepts and their relationships.

Example project plan

1. Learn about ontologies and try matching on small public datasets.
2. Build a program that uses clues from the ontologies to decide which matches to check first.
3. Test rankings against reference matches, then ask one or two LLMs about difficult cases.
4. Vary the number of checks. Compare match quality, conflicts, time etc.

Does wording change an LLM's answer?

Same meaning. Different answer?



Example project plan

1. Generate example questions in everyday language and formal notation.
2. Build a tool to test equivalent phrasings across several LLMs.
3. Compare answer accuracy, consistency across phrasings, resp. time and token use.
4. Test the best prompts in a tool that builds knowledge bases and check how reliable the answers are.

About me (Knut Omang)

- Ph.D from Ifi - high performance networks, “Il’er”@ifi since then
- Scali (→ Platform Computing): “Affordable supercomputing”
- Fast (→ Microsoft)
 - Search engines: My part: search within documents = “teasers/summaries”
- Paradiat (→ Logitech): Firewall traversal for realtime multimedia
- Oracle: Driver/kernel and Qemu (virtualization) development
- From Feb’20: CTO, Properate
 - Cloud based monitoring and control platform for commercial buildings
 - Monitors and controls commercial buildings - all over Norway, Sweden,Finland, Denmark++
 - Built on top of Cognite Data Fusion



Thesis areas

- Challenges related to data gathering and presentation of data from millions of sensors
 - feed from many types of sensor hardware
 - data modeling, ontologies, new ideas from data
 - creative new user interfaces/presentations (for instance: flexing)
- Development tools
 - Rust
 - Git reviewing – how will it be affected by AI
- Cloud/microservices related
 - Web services, REST APIs, protocols, functionality, security, IaaS
 - Virtual network technologies, containers, Kubernetes
- Using/exploring ideas using the programming language Rust
- Virtualization using QEMU
 - Innovations in QEMU, explore and understand features
 - Device emulation, passthrough/SR-iov, paravirtualization, nested virtualization
- Linux kernel development/test related topics
 - Kernel support for containers and VMs
 - Experiment with new APIs, options for building

