brikka.

MSc Thesis in AI Manufacturing Intelligence

Project: AIMI

Date: 2026-09-07

Thesis 1: Human AI Operating Models for Manufacturing Decision Support

Level: MSc thesis | Period: Spring 2027 | Reviewer: Didem Gürdür Broo | Supervisor: Erik Simu

Theme: Develop an advanced human-AI operating model for governing, validating, and using AI-supported decision systems in manufacturing.

Background: AIMI changes how operational and tactical decisions are prepared, validated, and owned. Manufacturing organizations therefore need clear decision rights, validation processes, escalation paths, and governance routines. The MSc thesis should examine how people and AI systems can work together under normal, uncertain, and exceptional operating conditions.

Suggested research questions:

- How should decision rights be divided between users, managers, system administrators, and the AI-supported system?
- What validation processes are needed before system outputs can influence production decisions?
- How should uncertainty, disagreement, and system failures be escalated?
- How can the operating model support long-term organizational adoption and improvement?

Suggested methods:

- Literature review on human-AI decision-making and AI governance
- Interviews or workshops with relevant stakeholders
- Decision-rights, workflow, and capability mapping
- Scenario-based evaluation of a proposed operating model

Expected outcome:
A human-AI operating model defining roles, decision rights, validation checkpoints, escalation paths, and governance routines. The result should provide practical recommendations for introducing and improving AI-supported decision-making in manufacturing.

Suitable student profile:
Students interested in industrial engineering, information systems, organizational design, operations management, AI governance, or human-AI collaboration.

Thesis 2: GraphRAG for Manufacturing Decision Support

Level: MSc thesis | Period: Spring 2027 | Reviewer: Didem Gürdür Broo | Supervisor: Erik Simu

Theme: Develop and evaluate graph-based retrieval and GraphRAG approaches for answering production questions with traceable evidence.

Background:
AIMI uses knowledge graphs to connect production processes, machines, materials, events, products, and decisions. At MSc level, this foundation can support more advanced retrieval and reasoning. The main challenge is to produce useful answers while preserving links to the underlying data and managing incomplete or inconsistent information.

Suggested research questions:
- Which graph representations and retrieval strategies best support selected production questions?
- How should graph context be provided to a large language model?
- How does GraphRAG compare with document-based or vector-based retrieval?
- How can the system preserve traceability and communicate data-quality limitations?

Suggested methods:
- Literature review on knowledge graphs, retrieval-augmented generation, and GraphRAG
- Definition of representative production questions and evaluation criteria
- Development of a graph-querying or GraphRAG prototype
- Comparative evaluation and error analysis

Expected outcome:
An evaluated GraphRAG prototype for manufacturing decision support, including documented design choices, comparison results, traceability mechanisms, and identified failure modes. The result should provide recommendations for AIMI's next technical implementation steps.

Suitable student profile:
Students interested in knowledge graphs, information retrieval, large language models, data engineering, computer science, or industrial analytics.

Thesis 3: Explainability and Validation for AI Supported Production Decisions

Level: MSc thesis | Period: Spring 2027 | Reviewer: Didem Gürdür Broo | Supervisor: Erik Simu

Theme: Develop and evaluate explainability and validation methods for trustworthy AI-supported production decisions.

Background:
AIMI is intended to support decisions that affect real production work. Users therefore need to understand the evidence, assumptions, and uncertainty behind system outputs. Explanations must support informed action while also enabling traceability, reproducibility, and responsible human oversight.

Suggested research questions:
- What explanations do different user groups need for operational and tactical decisions?
- How do different explanation formats affect understanding, trust, and decision quality?
- Which validation methods can identify unsupported or inconsistent outputs?
- How can explanations be linked to source data, graph relationships, and system versions?

Suggested methods:
- Literature review on explainable AI and decision-support evaluation
- Interviews or workshops with potential users
- Development of explanation and validation prototypes
- Comparative user evaluation using quantitative and qualitative feedback

Expected outcome:
A validated explainability and assurance framework for AIMI, supported by tested interface or interaction concepts. The result should define explanation requirements, validation procedures, traceability mechanisms, and recommendations for industrial use.

Suitable student profile:
Students interested in human-computer interaction, UX research, explainable AI, cognitive systems, industrial engineering, or applied AI.

Thesis 4: Specialized LLMs for Manufacturing

Level: MSc thesis | Period: Spring 2027 | Reviewer: Didem Gürdür Broo | Supervisor: Erik Simu

Theme: Compare methods for specializing large language models for manufacturing decision support in on-prem or hybrid environments.

Background:
AIMI aims to combine scalable architecture with domain-specialized AI for industrial environments. General language models can be adapted through structured context, retrieval, knowledge graphs, tools, or lightweight model adaptation. These approaches must be compared in terms of performance, traceability, resource requirements, and deployment constraints.

Suggested research questions:
- Which specialization approaches provide the best accuracy and groundedness for selected manufacturing tasks?
- How should retrieval, graph context, tools, and model adaptation be combined?
- What are the trade-offs between quality, latency, resource use, maintainability, and data protection?
- How well can the selected approach transfer across different tasks or use cases?

Suggested methods:
- Literature review on domain adaptation and industrial AI deployment
- Definition of representative manufacturing tasks and evaluation criteria
- Implementation and comparison of selected specialization approaches
- Benchmarking of output quality, latency, resource use, and traceability

Expected outcome
A recommended approach for specializing large language models for manufacturing decision support, supported by a prototype and comparative evaluation. The result should also include deployment considerations and an on-prem or hybrid reference architecture.

Suitable student profile:
Students interested in machine learning systems, large language models, software architecture, industrial AI, data platforms, or model evaluation.