

AI / OPTIMISATION DEVELOPER

LLM Agents + Scheduling | Hybrid | Team: IT | Full-time

We're growing our engineering team at heta. If you're excited by LLM agents and optimisation problems, enjoy measuring whether your work makes things better, and like owning what you ship, this role might be for you.

ABOUT THE ROLE

heta is building an AI-supported fiber network design, planning, and management platform. Two pillars of the product are genuinely AI-driven: an LLM agent system that reviews and validates network address data, and a planning-optimisation engine that schedules construction crews across thousands of work units under real-world constraints. This role lives at the heart of both.

You'll develop and improve these systems: tuning agent tools and prompts on one side and formulating & refining scheduling constraints on the other. You enjoy fastening AI reasoning to deterministic, production-grade code, and you validate quality on real data. Within our compact, fast-moving engineering team, you'll work closely with the full-stack engineers who expose this work to users.

WHAT YOU'LL DO

LLM AGENTS

- Build and refine LLM agent workflows: agent tools, structured outputs, multi-agent delegation, confidence scoring, and prompt engineering for domain-specific decisions.
- Integrate and manage language models (Anthropic Claude) in a production backend - model selection, token-cost tracking, and streaming responses.
- Keep AI suggestions conservative and reviewable: agent output maps to concrete, human-approved actions.

PLANNING & OPTIMISATION

- Develop the scheduling optimiser: constraint formulation, heuristic design, solver tuning, and warm-start strategies (Google OR-Tools CP-SAT, graph algorithms).
- Model real-world constraints - team skills, task dependencies, availability windows - and minimise project makespan.

QUALITY & OWNERSHIP

- Evaluate output systematically: design metrics, test on real datasets, and reduce error rates.
- Write clean, well-tested Python; participate in code reviews and design conversations; use Git fluently.

WHO YOU ARE

YOUR SKILLS

- Strong Python fundamentals, comfortable with typed, clean, well-structured code.
- Hands-on experience with at least one of: (a) LLM / AI engineering - building applications on top of LLM APIs (prompts, tool / function calling, structured output; Anthropic / OpenAI / similar); or (b) Optimisation / operations research - constraint programming, linear / integer programming, scheduling, or graph algorithms (OR-Tools, PuLP, scipy, networkx, or similar).
- Data fluency with pandas / numpy.
- A rigorous, measurement-driven mindset - you confirm a change improves results.

YOUR MINDSET

- You're motivated by getting AI and algorithms to produce reliable, trustworthy results on messy real-world data.
- You value ownership and take responsibility for what you put into production.
- You're curious, open to feedback, and enjoy open discussion of ideas.
- You appreciate a pragmatic engineering culture that balances experimentation with reliability.

NICE-TO-HAVES

- The other of the two pillars above - strength in both LLM engineering and optimisation is the ideal profile.
- pydantic-ai, Pydantic, FastAPI, async Python.
- Google OR-Tools (CP-SAT), scheduling / routing problems, makespan minimisation.
- Prompt engineering, agentic patterns, RAG / embeddings / vector stores.
- Geospatial data (geopandas, shapely, PostGIS).
- Workflow orchestration (Prefect), PostgreSQL, Docker, Azure.
- An academic or project background in operations research, applied mathematics, or machine learning.

WHAT WE OFFER

- A central role in the two most distinctive, AI-driven parts of the product.
- Real ownership of systems that make measurable decisions on real network data.
- A modern engineering environment that values collaboration, ownership, and open communication.
- Room to grow toward AI engineering, applied optimisation, or platform depth - supported by learning and hands-on experience.

HOW TO APPLY

Please send your application to HR@heta.eu with the title: **AI / OPTIMISATION DEVELOPER**