

Finland's Largest AI Research Group
in Software Engineering




GPT-Lab

From Code to Society



From Code to Society



Introduction

Founded in 2023 by Professor Pekka Abrahamsson, GPT-Lab has grown rapidly into the largest AI Research group in Finland. Focusing on the application of Generative AI in Software Engineering, GPT-Lab explores the impact of GenAI on development processes, tools and future worklife.

Vision



Future-proofing Software with AI



Mission

We conduct interdisciplinary AI research that transforms scientific advances into practical, real-world solutions for organizations and society.

We develop responsible, trustworthy, and increasingly autonomous AI that automates routine work and enables people to focus on more meaningful, creative, and high-value tasks.



8 RESEARCH PILLARS | 10M+ EUR FUNDING | IMPACTFUL, WORLD-CLASS AI RESEARCH



Projects


AI Champion

MAISA

ANSE

Evil-AI, BioAI, Syntetica, GENT, AICO, GETEK... 10+ more





**GPT-Lab operates in three campuses:
Pori, Tampere and Seinäjoki.**

**International collaboration in Brazil, Austria,
USA, Tanzania, Denmark, Sweden**

Services

For Companies

We focus on GenAI solutions that boost productivity, drive digital transformation, and build AI readiness

- Adopt an AI Researcher
- FAIR EDIH Activities

For Students

We provide

- hands-on courses
- thesis supervision
- real-world project opportunities

For Researchers

We focus on GenAI in Software Engineering, but welcome interdisciplinary collaboration in AI, software engineering, ethics

Impact

FUNDING (EUR)

10 M+

STUDENTS TRAINED
THROUGH COURSES

300+

RESEARCH PAPERS
RELEASED AND
PRESENTED AT TOP
ACADEMIC VENUES

20+

AI RESEARCHERS
ADOPTED TO
COMPANIES

5

DEMOS AND PILOTS
DESIGNED AND
TESTED IN REAL-LIFE

15+

Meet the teams!

GenAI in Software Engineering:
From Code to Society



Prof. Jyrki Nummenmaa

Natural Language
Foundations of
Generative AI



Sebastian Sonntag & Zheyng Zhang

Generative AI for Requirements
Engineering



Kai-Kristian Kemell

AI-Native Software
Practices



From Code to Society



Muhammad Waseem

AI for Software
Design and
Development



Jussi Rasku

Optimization and Industrial
Deployment of AI Systems



Mika Saari

Sustainable and Green
Generative AI



Meet Our Expert Team

From Code To Society



Salla Westerstrand

AI Governance



Sanni Pöntinen

AI, Work and Society



Case 1: Safe Backbone Routes for Unmanned Ground Vehicles

From Antti Pham's research
in the MixedFleet project
in collaboration with RovesUGV project


<https://www.fima.fi/mixedfleet>

Industrial Zone Autonomous Logistics: From shortest routes to consistent, safer backbone corridors

Observation: Trip-level optimization does not produce a coherent network-level structure.

Challenge: Trip-level behavior creates inconsistent routing patterns and safety risks in densely built areas.

Research: Discover backbone corridors and site access routes that provide a shared mobility structure for autonomous industrial logistics.



Industrial logistics needs bigger and heavier UGVs which makes them more dangerous.





Naïve (shortest) ✗:

Routes are everywhere; many segments used; uses smaller roads; road and crossing safety is not considered




Backbone-based ✓:

Concentrates on main corridors; considers road and crossing safety; fewer used segments; more consistent but longer trips





Identification of pathways through the industrial plots



Blue colors = backbone routes,
Green dot = pickup point,
Red lines = access routes;

The yellow area prevents the search for additional exits.

Routing approach (Dijkstra algorithm on the heatmap):
cost = distance + proximity to obstacles (safety distance).

Our contribution: From individual trips to a coordinated industrial mobility network

while supporting safety, AI-based planning, and autonomous UGV operations at scale.




Image source: <https://projektit.seamk.fi/rovesugv/> and Antti Pham

Democratizing Optimization for Smarter Cities

How AI agents can help uncover and solve optimization opportunities


Millenium Malamla

Research Assistant at GPT-Lab, Tampere University


Master's student in Software, Web and Cloud (AI-native Software)

Problem: Many Optimization Opportunities are Still Missed


Waste collection routes




Resource allocation



Public Transport Scheduling



Traffic Signal Timing



- Logistics and warehousing
- Maintenance planning

Why Are Optimization Opportunities Missed?

Common barriers

“We did not know this could be optimized.”

“We do not know which method to use.”

“We do not have optimization specialists.”

“A custom solution is too expensive.”

“There is not enough time to develop it

Operations research approach

→ Identify the optimization opportunity

→ Define decisions, objectives and constraints

→ Select an algorithm or heuristic

→ Build the software


→ Test and validate the solution

AI Agents Can Help Lower the Barrier to Optimization

At GPT-Lab, we are exploring how **multi-agent AI systems** could help turn real-world problems into working optimization experiments.



From a real-world problem to a better solution

Multi-agent AI for optimization



Case Example: AI-Assisted Warehouse Optimization observations

A fast warehouse simulator with an AI-assisted heuristic optimization loop, automatic KPI evaluation, and iterative strategy refinement.



Impact : Democratizing Optimization for Smarter Cities

If AI can lower the technical barrier to optimization, more organizations could experiment with optimization in areas such as:

- **Mobility** — traffic, routing, public transport
- **Energy** — buildings, district energy, demand response
- **Urban logistics** — deliveries, waste collection, warehousing
- **Infrastructure** — maintenance and asset scheduling
- **Public services** — resource allocation and emergency response



CONTACT DETAILS:

Millenium Malamla

Email:

millenium.malamla@tuni.fi

Linkedin: [Millenium Anthony |](#)

[LinkedIn](#)



RESEARCH PRESENTATION

An Agentic AI Factory That Learns to Optimize HVAC Control for Energy Savings

Rajratan Wankhade

Project Researcher, GPT-Lab (Tampere University)



Designing Smarter, Sustainable Cities Starts Indoors

Buildings sit at the centre of the urbanisation story and HVAC is where comfort, air quality and energy all meet.



WHERE WE LIVE OUR LIVES

- We spend roughly 90% of our time indoors (US EPA)
- Indoor pollutant levels can run 2–5× higher than outdoors
- Air handling units are what stand between people and that air




THE URBANISATION WAVE

- 45% of the world already lives in cities, a further 36% in towns (UN, 2025)
- Two-thirds of population growth to 2050 will happen in cities
- More people in more buildings means more air to condition and more energy to condition it

Koja × GPT-Lab: a real building, a real control problem


We partnered with Koja to optimize HVAC control for comfort and energy together, using telemetry from their own IoT sensors.



Agents propose, deterministic gates dispose

Agents interpret and plan. Scripts calculate every number and decide whether the run may continue.

LLM agent Deterministic gate Tool / script Human checkpoint



Nine agent roles do the interpreting. Not one of them grades its own output.



PASS

Continue to the next stage

The evidence supports the step, so the run proceeds — and the result stays linked to the exact data and settings behind it.



REVIEW

Bounded replan, then retry

The planner revises and comes back through the same gate. Capped at two attempts, so it cannot loop forever.



REFUSE

Terminal halt — escalated to a human

Missing data or a site meaning nobody has confirmed will not improve on retry, so the run stops and waits for a person.

LangGraph keeps the shared state, routes work between nodes, pauses for review, and preserves every run's history.

Learning to control a building without touching it

One model learns what the building does. The other learns what to do about it — by practising inside the first.

1 - Surrogate model: a learned copy of each zone



16/16 one-hour and 80/80 rollout comparisons beat persistence

2 - Policy optimization: the policy practises inside the surrogate



The action space is mixed — a discrete operating mode plus a continuous airflow setpoint — which is what makes this control problem hard.



STEP 1 · SURROGATE

A learned copy of each zone

It reads the last 48 measured hours and predicts what happens next — tracking that a room has been warming for six hours, not just the reading right now.



STEP 2 · POLICY

Practises inside the surrogate

The controller chooses operating mode and supply airflow, scored on comfort plus air quality minus fan use. It never touches the live building.



Everything here is offline

Saying so plainly is not a disclaimer — it is the difference between a claim I can defend and one I cannot.

What this would mean for Koja’s three zones

Less fan energy and better comfort in every zone — measured offline against what the building actually did.



FAN ENERGY SAVED

10.9%

across the whole site, weighted by each zone’s baseline




HOURS IN COMFORT BAND

86.4% → 96.6%

+10.2 points, across equal zone-hours

Where the fan energy goes today



■ TK1 · Production 94% ■ TK2 · Offices 5% ■ TK3 · Gym 1%

TK1 carries almost all of the site’s fan energy, so its small % is the biggest absolute saving.

TK1 · PRODUCTION

Comfort held at 100%

The biggest absolute energy saving on site

–7.8%

fan energy

TK2 · OFFICES

Comfort 87.3% → 94.5%

Largest proportional saving — watch CO₂ here

–62.7%

fan energy

TK3 · GYM

Comfort 72.1% → 95.3%

The biggest comfort turnaround of the three

–19.5%


fan energy



The trade-off we are not hiding: mean CO₂ rose from 441 to 486 ppm, and TK2 crossed 800 ppm for 3 of 14,400 simulated hours. These are offline LSTM-simulated results — the policy has never controlled the live building.

What this demonstrates —> and the path from here

What we proved, what it delivered, what to fix, and how it reaches the next building.



From one-size-fits-all operation to an optimized solution for every zone : without removing human responsibility.

Johannes Määttä

- Research assistant at GPT-Lab
- Master's student in Human-Technology Interaction at Tampere University
- Presenting TARMO:
 - TARMO was created as part of the AINA research project.
 - TARMO is an LLM-powered newsdesk designed for local newspapers.

<https://www.linkedin.com/posts/aina-tampere-pori-ugcPost-7475479959885705216-22lv/>

<https://www.linkedin.com/in/johannes-maatta/>





www.gpt-lab.eu

Follow us for more!



**5000+
followers**