

Senior Machine Learning Engineer

Full time -Toronto

At MapsPeople, we are entering a pivotal phase of AI-led transformation and growth, shaping the future of spatial intelligence and indoor navigation for large enterprise customers worldwide. To support this evolution, we are looking for a Machine Learning Engineer to take ownership of the systems that power our indoor mapping platform.

The Role

You will own the full lifecycle of our machine learning systems from data ingestion and model training to deployment, optimization, and production monitoring. This is a hands-on role at the intersection of **computer vision, geospatial data, and real-world engineering**, where model performance directly impacts how people navigate complex indoor environments.

You will work on challenging problems such as generalizing models across diverse architectural inputs, transforming raw predictions into usable spatial data, and optimizing inference at scale. As part of a lean, high-impact team, you will have the opportunity to shape the direction of our ML pipeline, explore next-generation approaches, and deliver solutions that are directly reflected in the product experience.

We imagine you come with

- 5-8 years of experience building **machine learning systems in production**
- Strong proficiency in **Python**, with clean, maintainable, and well-tested code
- Experience with **deep learning frameworks** (PyTorch preferred)
- Strong background in **computer vision** (object detection, segmentation, or classification)
- Understanding of **ML systems and infrastructure**, training pipelines, dataset management, and model versioning
- Experience working with **cloud platforms** (GCP, AWS, or Azure)
- Familiarity with **containers, orchestration, and production deployment**
- Experience with end-to-end production pipelines including data ingestion, preprocessing, model training, inference, and monitoring.
- Strong understanding of **model evaluation and performance metrics**
- Ability to think beyond “model accuracy” and evaluate real-world impact

You should also have:

- Experience with geospatial data (coordinate systems, geometry processing, spatial databases)
- Familiarity with architectural drawings or structured building data (CAD, BIM, etc.)
- Exposure to foundation models or vision-language models
- Experience optimizing models for inference (quantization, runtime optimization, hardware acceleration)
- Background in adjacent domains such as robotics, autonomous systems, medical imaging, satellite imagery, or document AI
- Engagement with the ML community through open-source contributions, competitions, or research
- A degree in Computer Science, Machine Learning, Mathematics, Physics, Engineering, or a related field is one path
- But strong candidates also come from bootcamps, self-directed learning, adjacent

disciplines, or unconventional career paths

What you get

- Problems Worth Solving:**

Your work directly impacts how people navigate real buildings, from hospitals to airports. The feedback loop is tangible and immediate.

- Technical Ownership:**

You will own critical parts of the ML pipeline in a lean team where your decisions shape architecture and direction.

- Room to Explore:**

We are actively adopting modern ML approaches — you won't be maintaining legacy systems.

- A Unique Domain:**

Indoor mapping sits at the intersection of computer vision, geospatial engineering, and product thinking — a niche with real depth and complexity.

- Global Collaboration:**

Work in an international environment with teams across multiple regions and strong Danish roots.

- Flexible Work Setup:**

This role is based in the Toronto / GTA area with a hybrid work model.

- Competitive Compensation:**

We offer compensation and benefits aligned with the Toronto market and your experience level.

Ready to map the future with us? We'd love to hear from you - Apply now!