SILO AI

Building blocks and manual for the Edge AI of tomorrow

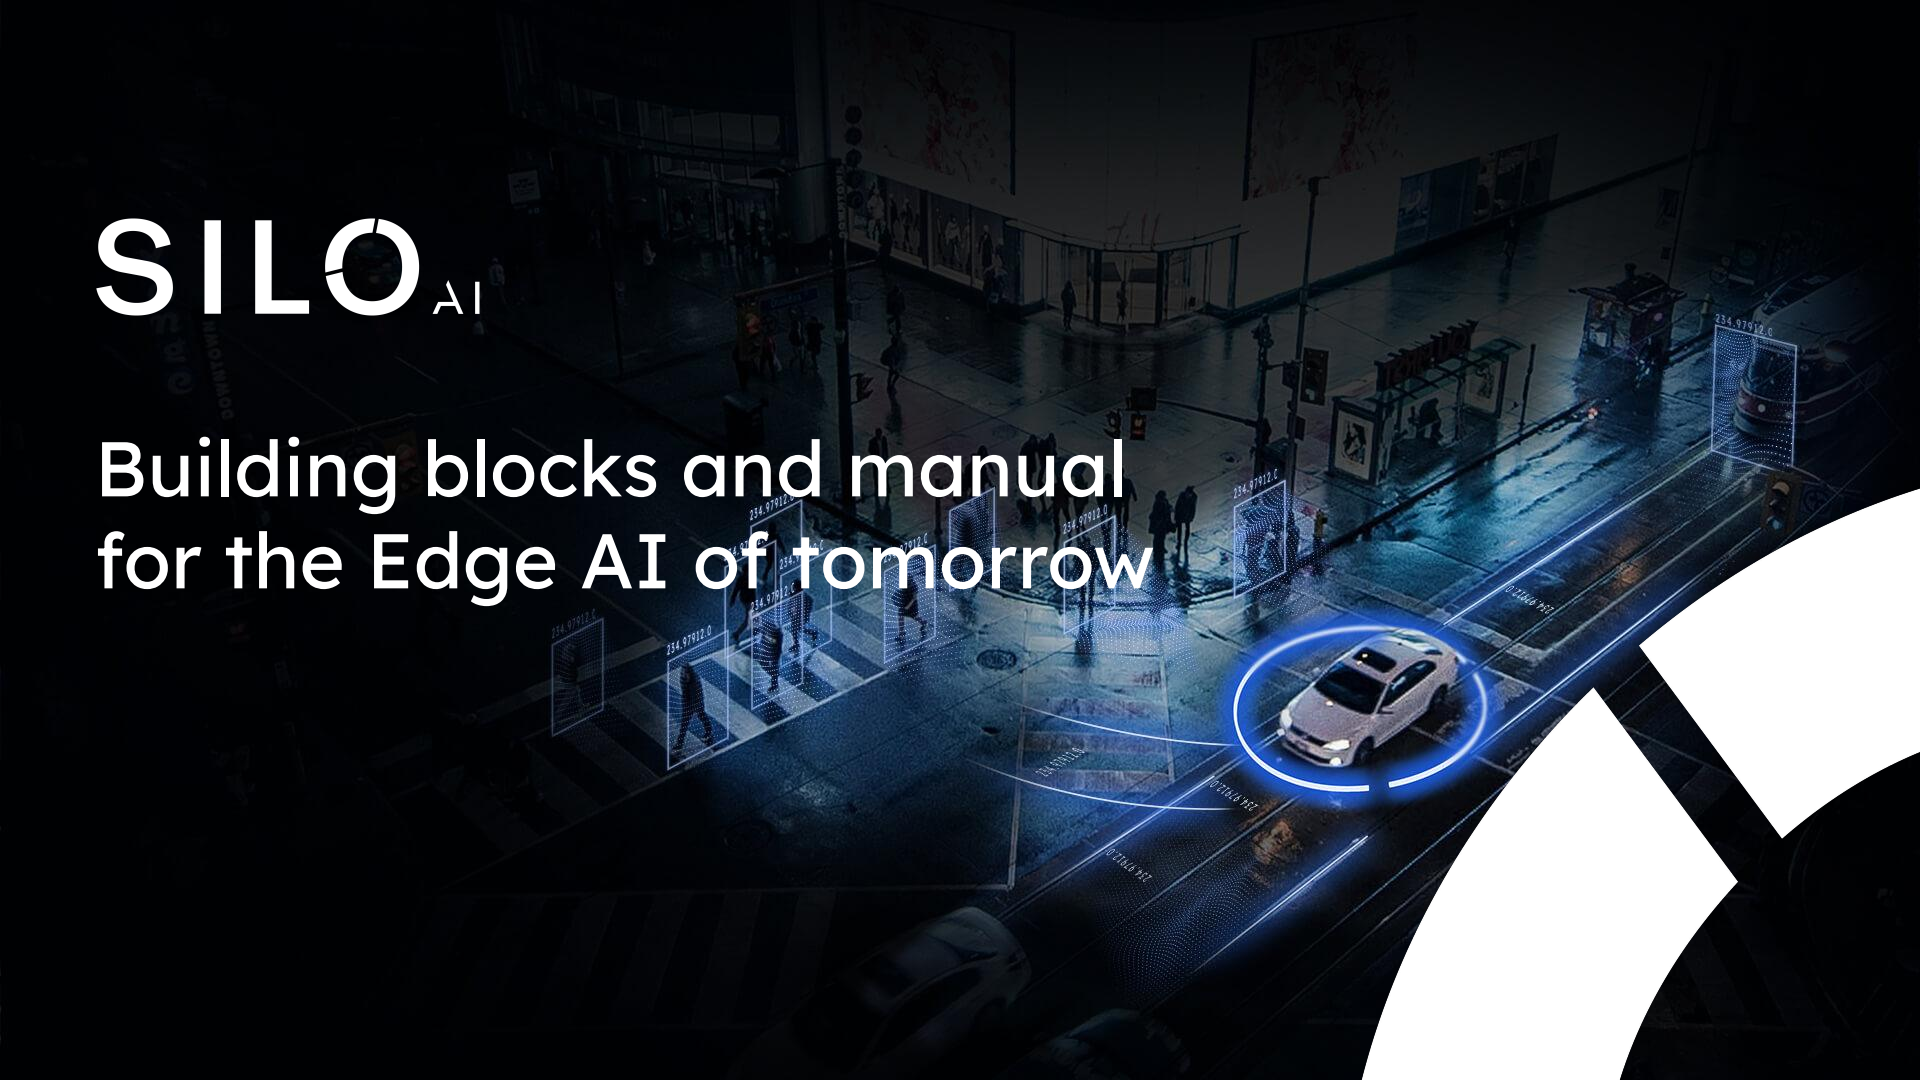

Who we are

Largest private AI lab in the Nordics

SILO AI

What we do

Trusted AI partner. We deliver AI-driven solutions and products to our clients by providing world-class expertise and tooling.

Vision

AI for people. A world with safe human-centric AI that frees the human mind for meaningful work.

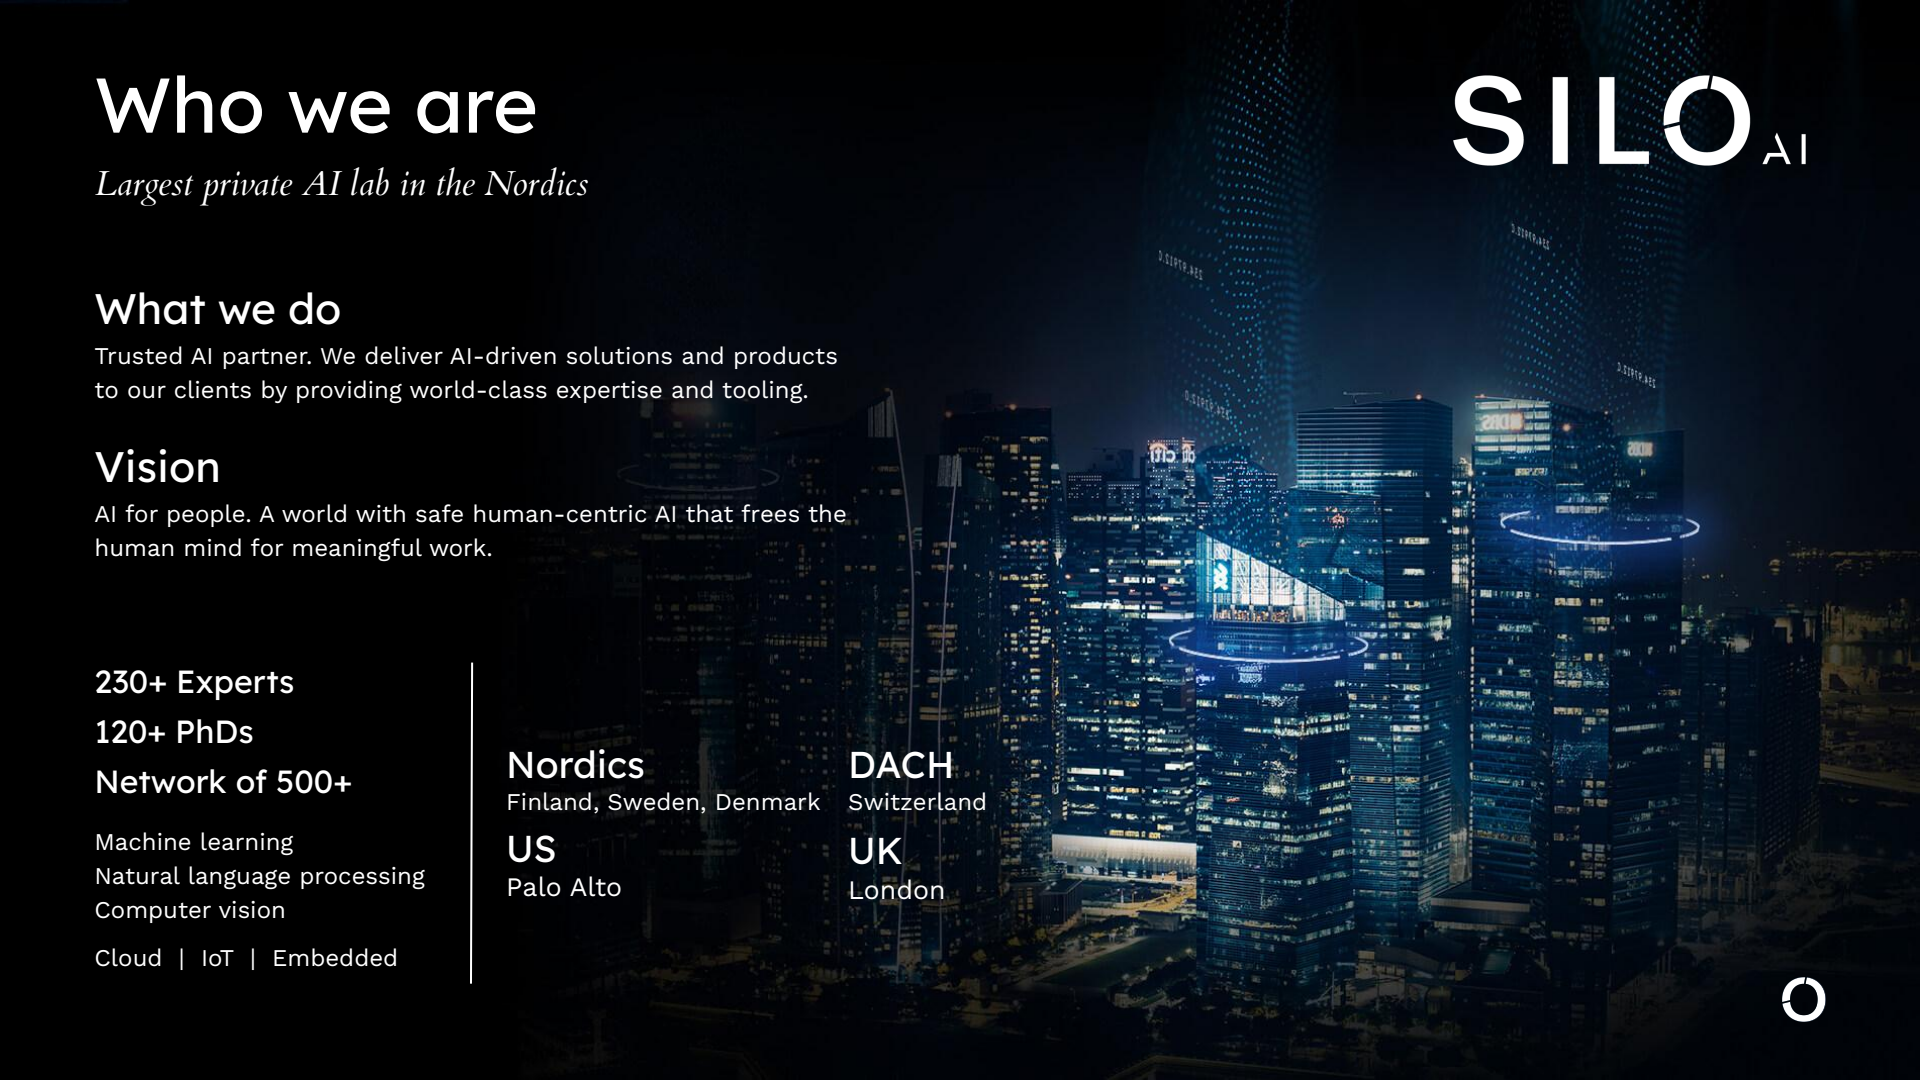

230+ Experts

120+ PhDs

Network of 500+

Machine learning
Natural language processing
Computer vision

Cloud | IoT | Embedded

Nordics
Finland, Sweden, Denmark

DACH
Switzerland

US
Palo Alto

UK
London

Silo AI Universe

Silo AI is building a world with safe human-centric AI. We build AI for people.

Silo AI was founded in 2017 to connect leading AI scientists with real-world problems and help companies succeed in building cutting-edge AI.

In short, we set out to be a trusted AI partner, to assure that companies stay competitive at a time when AI is globally being widely adopted. We set out to make sure that Europe has a flagship AI company.

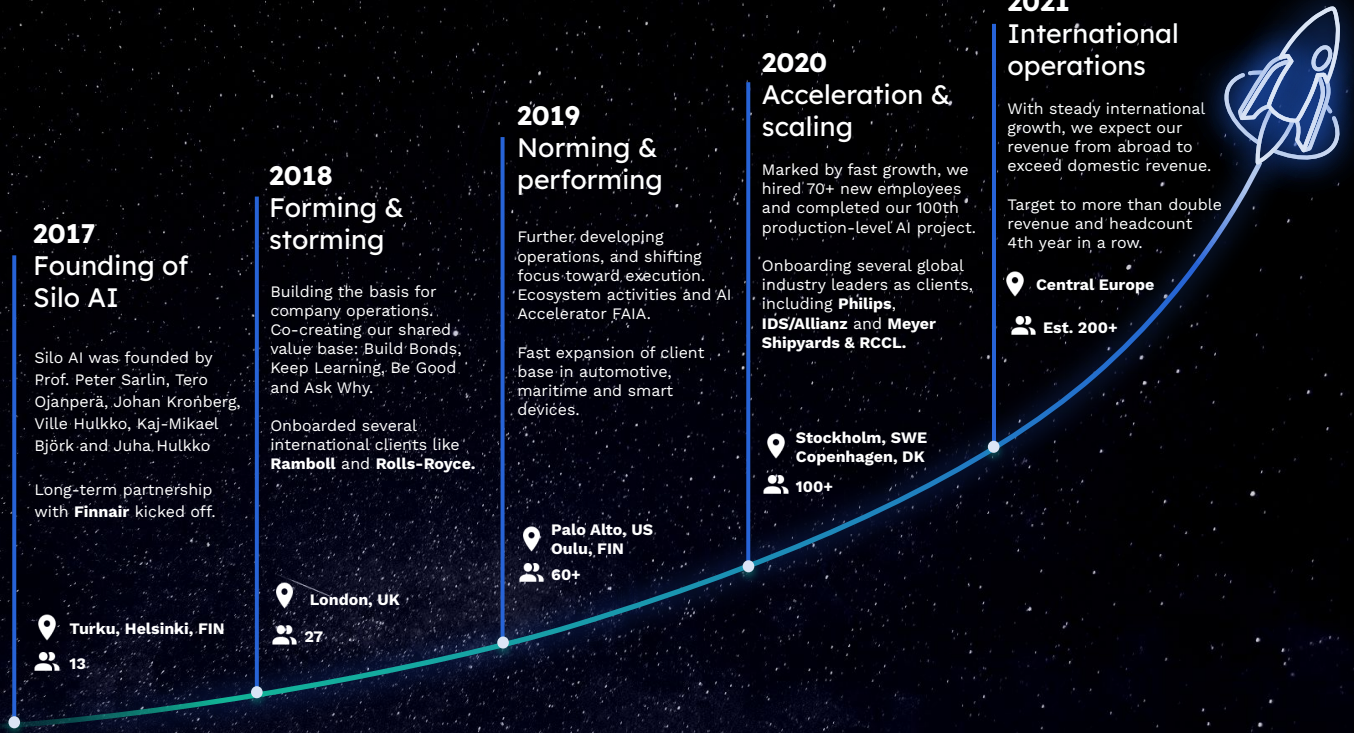

Why Silo AI?

State-of-the-art AI expertise

Deep and specialized AI expertise, a unique ability to attract and keep world-class AI talent

Trusted AI partner

Trusted advisor with a service mindset and flexible working model for a unique customer experience

End-to-end capability

With expertise spanning the tech stack and operations, realize best outcome and maximum value

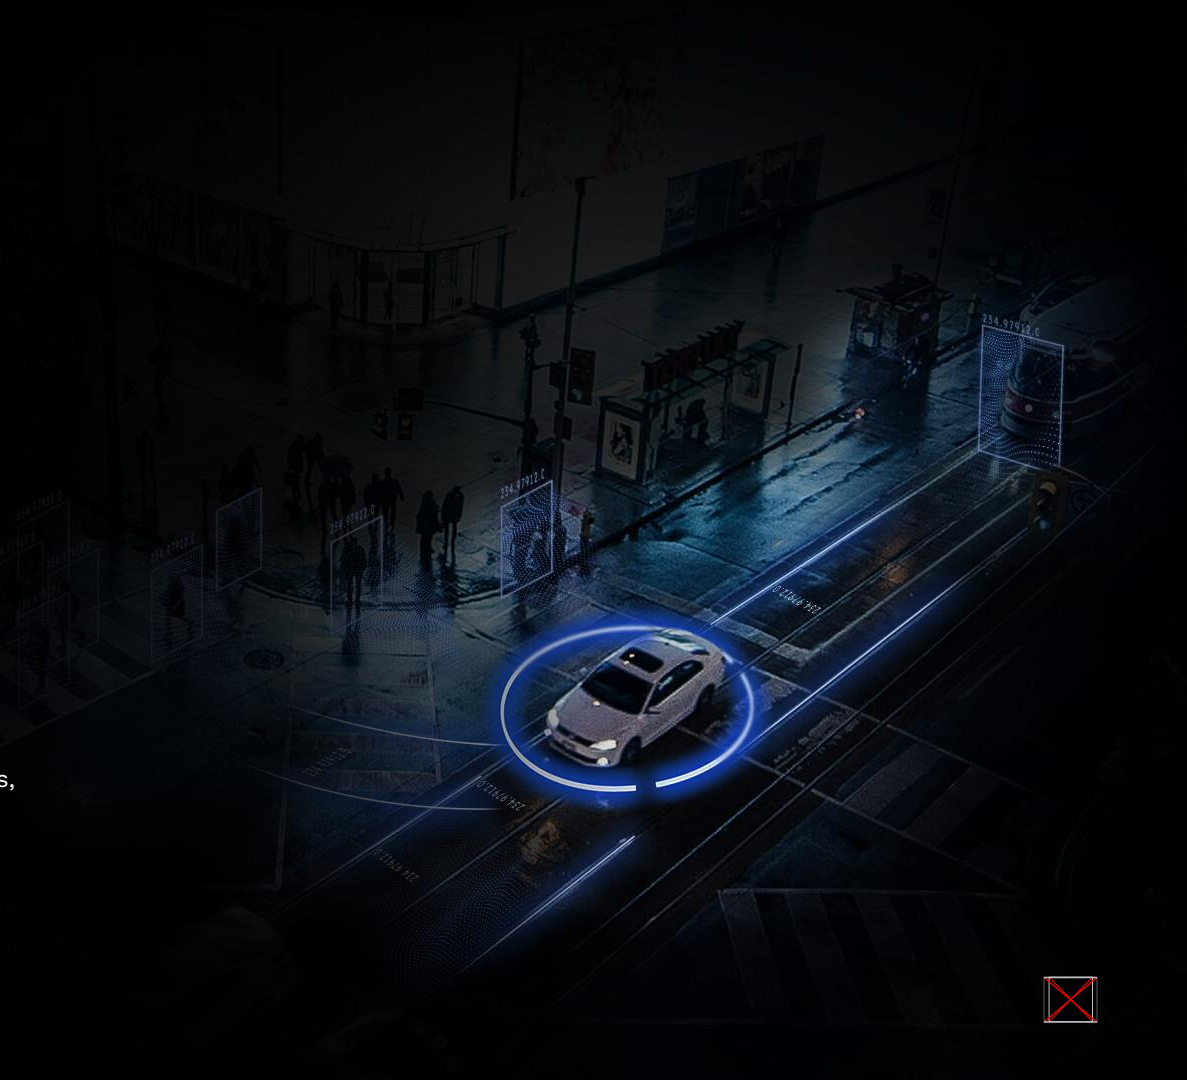

R&D for AI-driven products

Expertise areas

Design, planning & management

- Technology planning
- Architecture planning
- Rapid prototyping
- Project management

Machine Learning

- Supervised learning
- Unsupervised learning
- Transfer & active learning
- Multi-objective optimization
- Reinforcement learning
- Recommender systems

Computer Vision

- Image segmentation
- Object detection
- Sensor fusion
- Image generation
- 3D, video, point clouds
- Image & object annotation

Natural Language Processing

- Word & doc classification
- Text tagging & parsing
- Language generation
- Machine translation
- Speech recognition
- Conversational AI

Software, Data & ML Engineering

- MLOps & DevOps
- Data engineering
- Software development
- Lifecycle support

Data platforms

Cloud AI

Embedded AI

IoT & Edge AI

Use cases and references

Smart vehicles

Automotive | Maritime | Mining | Forestry | Aerospace | Heavy machinery

- Advanced Driver-Assistance Systems
- Intelligent Vessel Awareness System
- Autonomous Heavy Machinery
- Predictive Aircraft Maintenance
- Situational Awareness & Object Tracking

Smart industry

Machinery | Manufacturing | Process industry | Industrial IoT | Chemicals | Cleantech

- Visual Quality Control
- Visual Anomaly Detection
- Predictive Maintenance
- Process Quality Prediction
- Tool & Equipment Recommendation

Smart devices

Consumer electronics | Medical Devices | Networks | Semiconductors | Security

- Camera Device For Surveillance
- Cybersecurity & Intrusion Detection
- Sewage Pipe Defect Analysis Robot
- On-device Point Of Interest Prediction
- Healthcare & Medical Devices

Smart health

Medical Devices | Pharma | Healthcare Services | Home Care

- Surgery Imaging Analysis
- Imaging For Cancer Diagnostics
- Biomedical Signal Processing
- Treatment Recommendation
- Computer Vision For Home Care

Smart society

Energy | Logistics | Infrastructure | Retail | Construction | Mobility | Finance

- Water Quality & Sewage Pipe Analysis
- Port Automation
- Parcel Arrival Prediction
- Energy & Distributed Heating Optimization
- Financial Stability Monitoring

Edge AI

The TLDR

1) *Edge AI isn't the future, it's already the present*
2) *Edge AI isn't the model, it's the whole stack*

Why Edge AI?

Edge is solving a real problem:

"Why can't we just run it in cloud?"

Too illegal
Privacy laws,
private contracts

Too expensive
Huge data from a tiny unit
in a hard-to-reach place

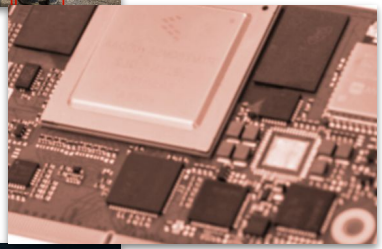

Too unreliable
Safety or business critical
operations

Too impossible
Air-gapped systems

What is Edge AI?

Among the explosion, a core structure remains.

Similar questions, different choices across use cases.

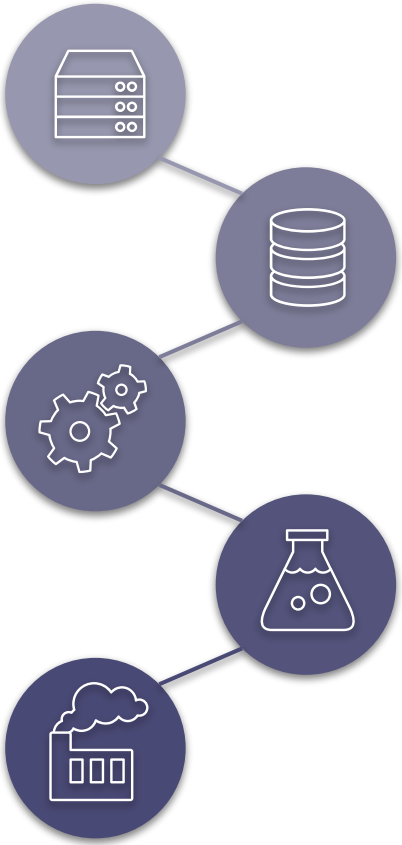

Two examples of the structure

	Emulsion inference	Industrial camera design
Hardware	Local mini-cluster, QoS for data flow, models, apps	NXP SoC + to-be-evaluated AI acceleration FPGA
Data	Reliable pipeline for messy field setups, secure data sharing	AI compute placement, AI on/off-flow image processing
Software	Deployment-specific customization of models and apps	Open AI-enabling middleware + TF lite
AI modeling	CNNs for inverse problems & auto-calibration	Latency and memory optimization, neural feature primitives
Operations	Operator support, rapid deployment tooling	Long-term hardware commitment to support future software upgrades

The end-to-end of AI

Building AI with the whole structure in mind is an essence, not an excuse

+ props to Clayton Christensen

AI is the resource-hungry, least mature piece of the average technology stack

⇒ Getting most out of state-of-the-art AI models requires concerted effort from all parts of the stack

⇒ Each part of the stack needs AI expertise to succeed

Edge AI is a living proof of this

How do AI and Compute interact?

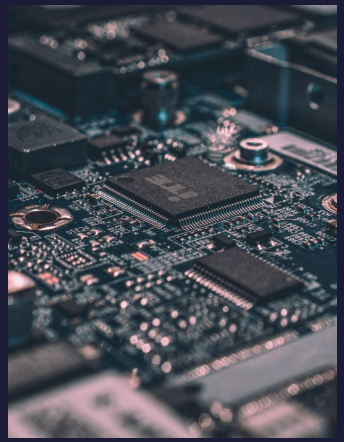

Environmental sensor
Joint design of MCU HW and ML model to a power budget

Wide area video analytics
Tune models for compute & memory profile and minimize “device tree” cost

Machine portfolio
Scalable solution framework across various sensor packages and AI use cases

How do AI and Data interact?

Automotive autonomy
Continuous ML-based sensor calibration

Quality control
Sensor setup design, data selection for edge vs cloud AI use cases

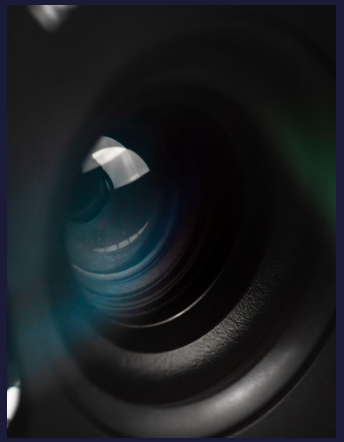

Smart device imaging stack
Share work across ISP and AI, fit to specific hardware

How do AI and Operations interact?

Machinery control
Safety and quality in predictive control for high power machinery

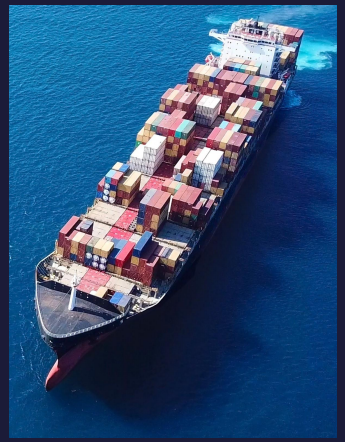

Vessel autonomy
Vessel<->cloud data selection, updates, and compute placement

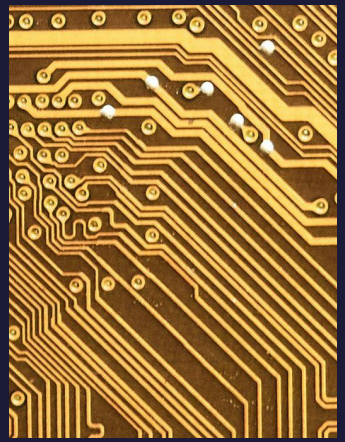

PCB scanning
Rapid deployment of new test stations with semi-automatic model tuning

Why MLOps?

Processes, systems, and infrastructure, but mainly mindsets

MLOps, as of now, is getting built as a cloud-centric methodology

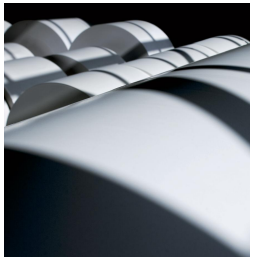

Rapid iterative learning
Learning to test, develop, and deploy new AI initiatives requires cheap repetition

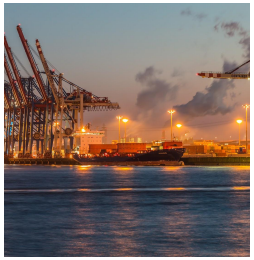

Systematic treatment of complexity
MLOps is like DevOps, but for data logic

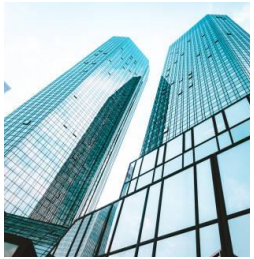

Data-centric mindset
Common self-reported AI challenges: data quality, corner cases, data drift

So what is MLOps for the edge ?

Richness of use cases still limits the formation of a common framework

Solution lies, for now, in picking and choosing from proven tactics per context

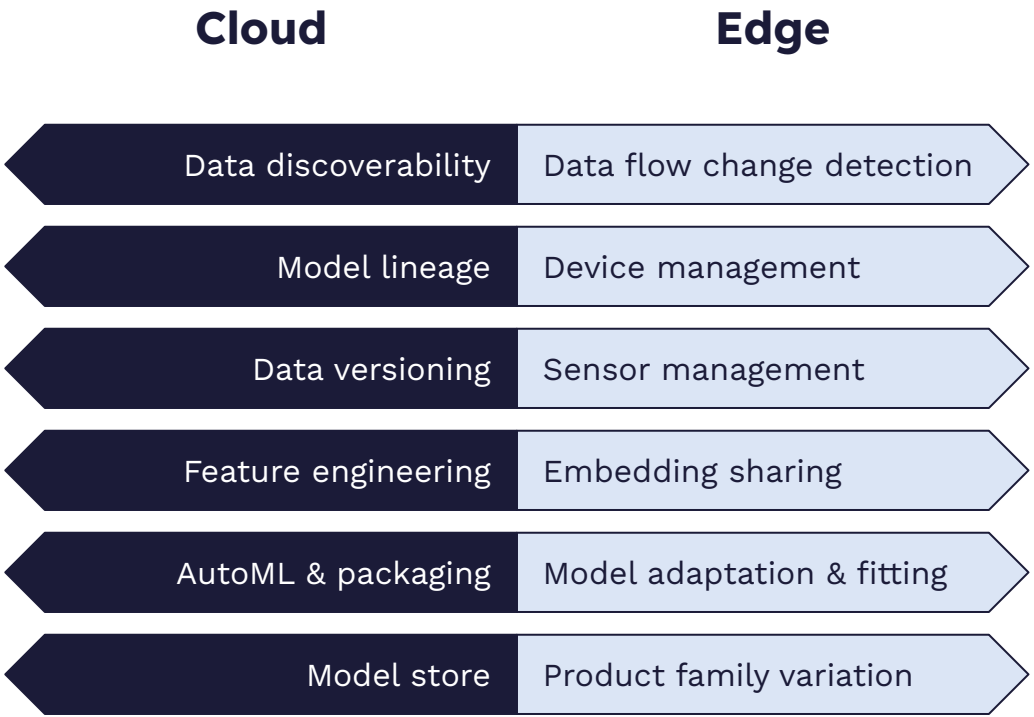

Edge AI

The TLDR

It's already a big industry

Common underlying structure

End-to-end design

Niko Vuokko

Head of Technology
PhD

niko.vuokko@silo.ai
+358 50 485 0513
https://silo.ai

Largest Private AI Lab in the Nordics

SILO AI