

SILO_{AI}

Largest private
AI lab in the Nordics



Who we are

Largest private AI lab in the Nordics

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.

180+ Experts

90+ PhDs

Network of 500+

SILO_{AI}

Nordics

Helsinki, Turku, Tampere, Jyväskylä, Oulu
Stockholm
Copenhagen

UK

London

US

Palo Alto




Our expertise




Machine learning

- Understand, decide and predict in highly complex systems
- Provide human experts with machine-driven support
- Extract value out of vast, complex combinations of data



Computer vision

- Automated image/video analysis for insights and suggested actions
- Machine vision for autonomous vehicles and robot guidance
- Augment human intelligence through video-based AI solutions



Natural language processing

- Extract information from docs, speech, news, articles, feeds, etc.
- Augment RPA applications with learning language processing
- Remove clerical bottlenecks, augment and empower experts



Solution design & development

- Development and design for turnkey solution delivery
- Software and data engineering for end-to-end development
- Integrate AI seamlessly into solutions and products




Covering All Streams



Models, algorithms and neural networks

Bring in specific expertise to strengthen existing teams and initiatives with deep, complementary skills in the AI from development to deployment



Embedded, edge and cloud

Analyse and plan the deployment and computing options to ensure that the solution also work in the real world and the tools and frameworks meet requirements



Enable efficiency with automation and tools

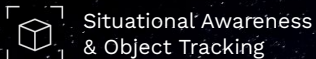
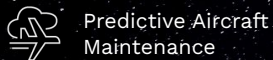
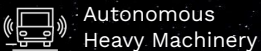
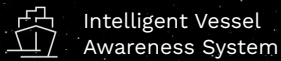
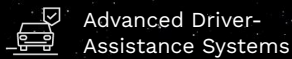
Provide team and best practises helping to design ML Ops architecture and tool set that scales cost-efficiently considering security and compliance and bringing automation to increase the efficiency and ease-of-work for data scientist



100+ production-level AI projects


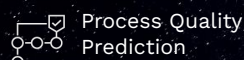
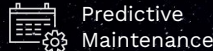


Smart vehicles

Automotive | Maritime |
Heavy Machinery | Aerospace




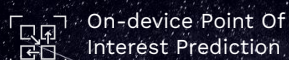

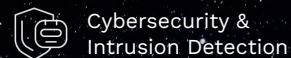
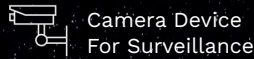
Smart industry

Manufacturing | Factories |
Process Industry | Industrial IoT



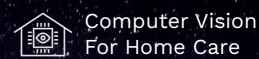
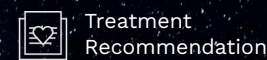
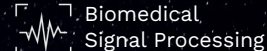


Smart devices

Wearables | Smart Homes |
Personal Electronics | Semiconductors



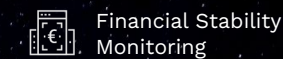

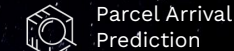


Smart health

Medical Devices | Pharma |
Healthcare Services | Home Care



Smart cities

Infrastructure | Energy | Logistics |
Construction | Mobility | Citizen



Optimized model
deployment on
edge/embedded



ML in low resource HW

Implementing a Machine Learning solution is not different from any other product implementation, it requires:

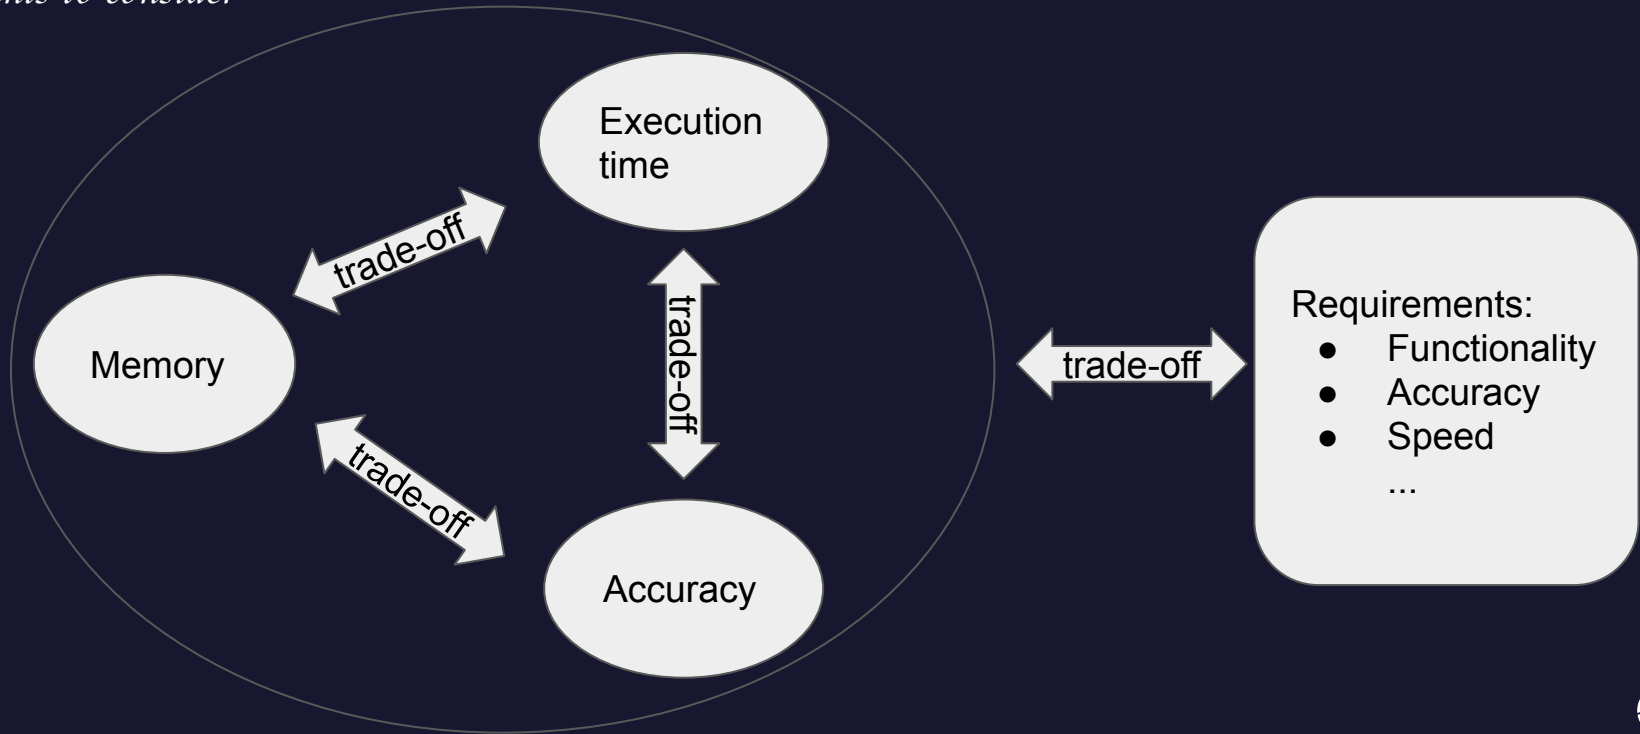
- Understanding the use case requirements.
 - Functionality, accuracy, speed, latency, maintainability, interfacing, ...
- Identifying system limitations.
 - Processing power, memory size, thermal limits, SW capabilities ...
- Taking into account R&D constraints
 - Cost, time-to-market, resources, available competences, ...

It is important to start with correct (rough) understanding of platform capabilities vs. ML solution complexity. Trying to modify too big ML model into **limited computing environment** or **restricted APIs** is often unsuccessful or time consuming use of R&D resources.



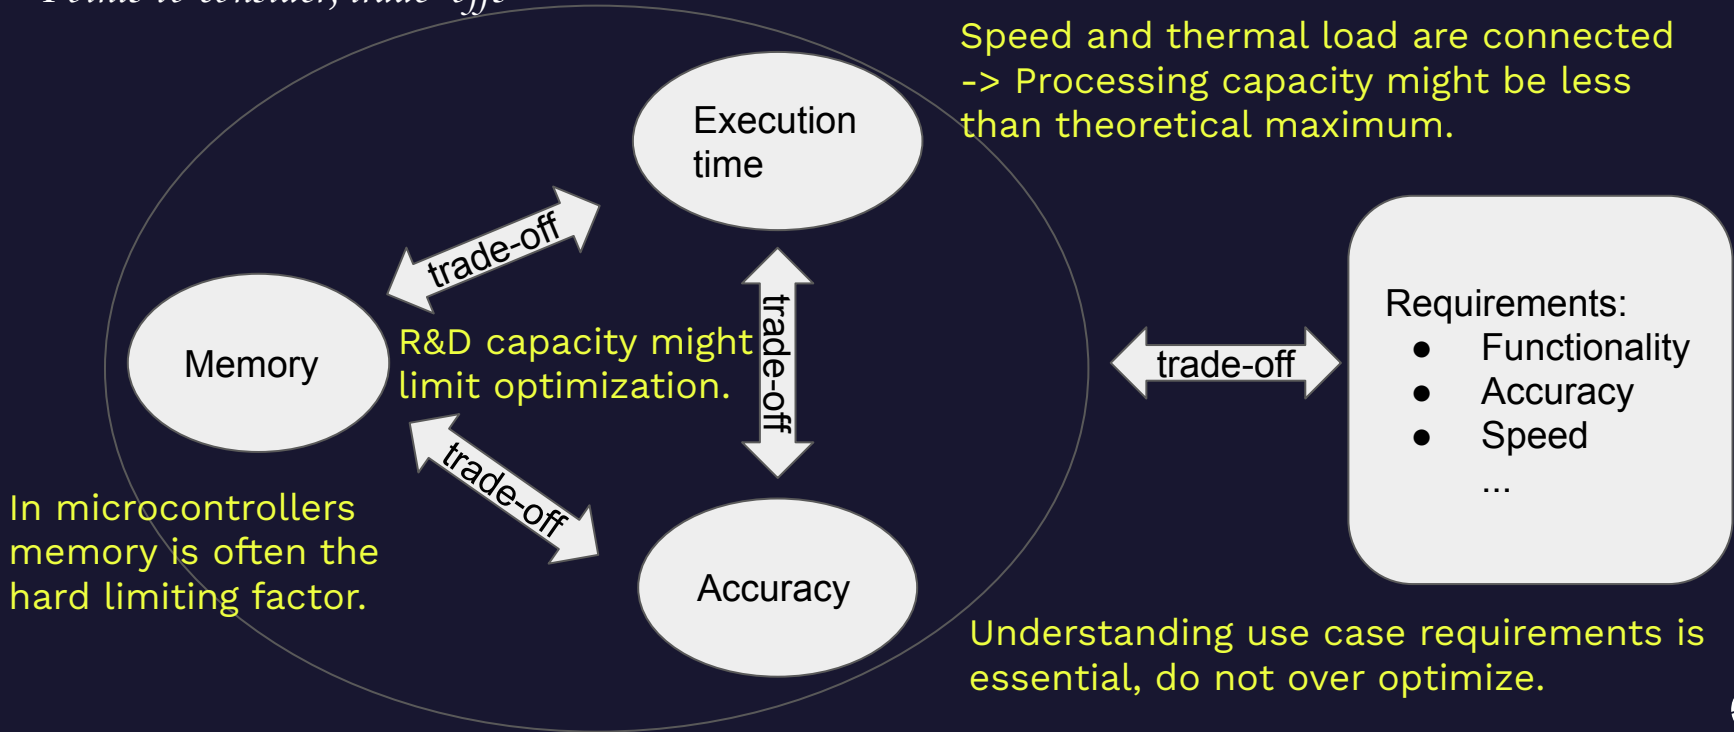
ML in low resource HW

Points to consider



ML in low resource HW

Points to consider, trade-offs



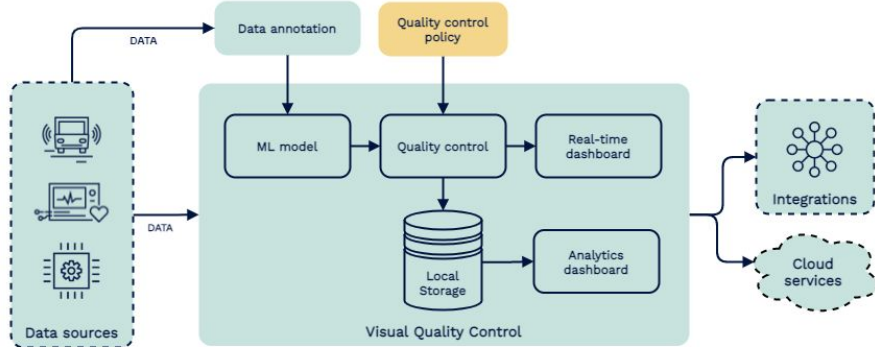


Results on Intel optimization

Silo OS Visual Quality Control

A modular and scalable solution for industrial use cases

- Test bed for developing end-to-end industrial quality control solutions
- Customized based on client needs (sensors, model type, integrations, ...)
- Optimized deployment to
 - Reduce inference latency
 - Improve throughput (multiple sensors)
 - Use available compute to the fullest even when sharing resources



Overview of optimized deployment process



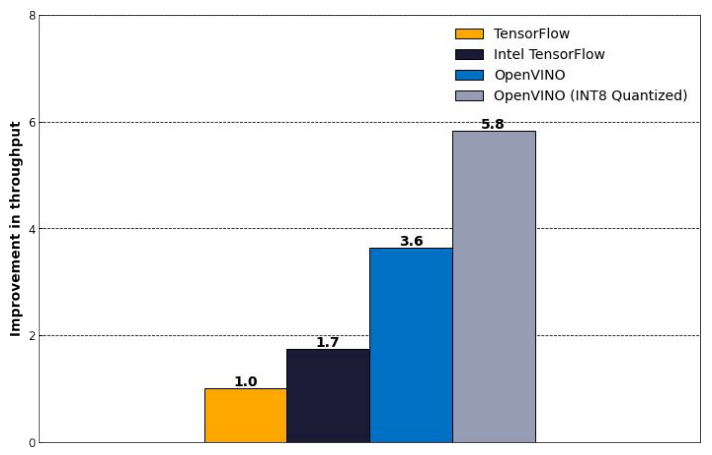
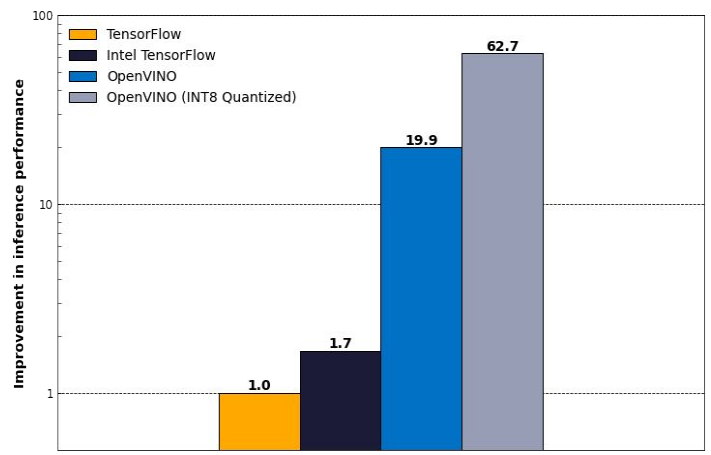
- Model trained with standard deep learning libraries, deployed in OpenVINO format to maximize performance on Intel CPUs
 - Supports TensorFlow (1.X and 2.X), PyTorch via ONNX and other formats
 - Supports common architectures, extendable with custom operations

- Exported models can leverage all Intel hardware: CPUs from Atom to Xeon, integrated GPUs, edge VPUs, FPGAs



Optimization results

Multiclass surface defect segmentation



- Significant gains in inference speed (ms/image) and throughput (images/s)
- Consistent results on different platforms and using different network topologies (segmentation, classification, ...)




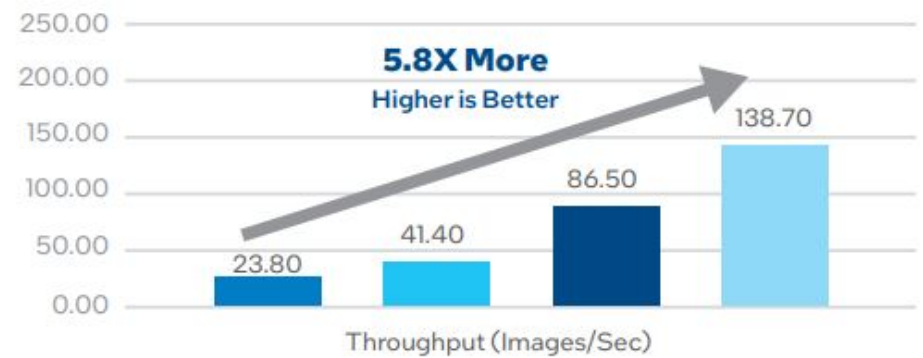
Optimizations Deliver Up To 62.7X Faster Processing¹

Silo AI completed benchmarking on Intel DevCloud using multiple architectures. Benefits using the Intel Distribution for OpenVINO toolkit include:

- Up to 3.8X more FP32 throughput on 3rd Gen Intel Xeon Scalable processors.
- Up to 3.1X more FP32 throughput on Intel Movidius Myriad X VPU.
- Up to 6.1X more INT8 throughput on 3rd Gen Intel Xeon Scalable processors.

Optimizations reveal up to 5.8X more images per second can be inferred with up to 62.7X lower latency using OpenVINO deployment compared to open source TensorFlow on Intel Xeon Platinum 8360Y processors.

Inference Throughput and Latency on Intel Xeon Platinum 8360Y Processor



■ Unoptimized Open Source TensorFlow ■ Intel Optimizations for TensorFlow ■ OpenVINO (FP32) ■ OpenVINO INT8 Quantized



The background is a dark blue, high-tech circuit board with intricate patterns and glowing points of light. In the center, a square chip is highlighted with a bright blue glow. To the right, a large, abstract yellow shape, resembling a stylized 'C' or a partial circle, is set against a black background with small white dots.

Optimizing for embedded/edge systems

Examples on NVidia



Intelligent awareness system for autonomous vessels

An AI solution to produce full awareness of the vessel's surroundings that operates with different sets of sensor data and combines vessel's own sensor fusion data with radio signals.

The awareness and guidance module of the awareness system tracks activities of other vessels at sea and provides information and recommendations for the crew.

The system also includes an integrated neural logic model for dynamic sensor packages and configurable end-user use cases.

Our role was to design and implement the solution architecture including an onboard EDGE solution and cloud integration.

**AUTONOMOUS
VESSELS**

INDUSTRY

Maritimes

CLIENT

Company developing
autonomous vessels

TECHNOLOGIES

Computer vision,
machine learning,
sensor fusion

DELIVERABLES

Platform architecture and development
Tailored annotation tool to speed-up labelling

SILO OS

Annotation platform





Watch the webinar recording from June 2, 2021

Sensor fusion for situational awareness



Recorded live on Wednesday June 2, 2021

Silo AI x NVIDIA webinar – Vessel and vehicle awareness, Case Groke Technologies

*This webinar is for you if you're an R&D lead, Director of AI development, CDO/CXO, or
Business/AI lead working with AI-driven sensor fusion.*

In this webinar, we dived into using sensor fusion to improve situational awareness. Together with NVIDIA Inception, we shared key use cases enabled by NVIDIA components, and we showcased invited guest Groke's concrete use of sensor fusion to improve vessel awareness.

Brought to you by

SILO_{AI}



INCEPTION
PROGRAM





Kai Knuutila

Lead AI Solutions Strategist

kai.knuutila@silو.ai
+358 50 351 5113
<https://silو.ai>



Largest Private AI Lab in the Nordics

SILO_{AI}