Q.ANT

Q.ANT NATIVE PROCESSING UNIT GEN 2

THE PHOTONIC AI ACCELERATOR

For energy-efficient and fast HPC and AI

THE FUTURE OF AI & HPC IS LIGHT

As AI scales globally, the traditional chip technology reaches its limits. Unlocking the true potential of AI demands a fundamental rethink of computing itself.

Q.ANT Native Computing ushers in a new era — computing with light instead of electricity promising 30x energy savings. The Q.ANT Native Processing Unit (NPU) is the world's first commercially available photonic analog processor computing with light. By combining high-precision photonic processors with advanced algorithms based on non-linear math, Q.ANT Native Computing will extend today's compute architecture and enable next gen AI and HPC applications in machine learning, computer vision and physics simulations. The future of AI and HPC isn't silicon. It's light.

OVERCOMING TODAY'S COMPUTE BOTTLENECKS WITH LIGHT

Q.ANT Native Computing tackles three fundamental challenges of today's compute ecosystem:

The progress in CMOS slows down
The miniaturization of transistors reaches its physical limits. State-of-the-art CMOS processors face performance shortcomings driven by an exponential growth in AI workloads.[1]

Energy consumption of AI data centers surges
Data center electricity consumption is set to more than double by 2030[2] limiting the upscaling of AI applications.

Limited semiconductor production capacity
AI hardware production is centralized in few regions and dominated by a small number of companies creating technological dependence.[3]

Semiconductor AI Processor Performance[4]

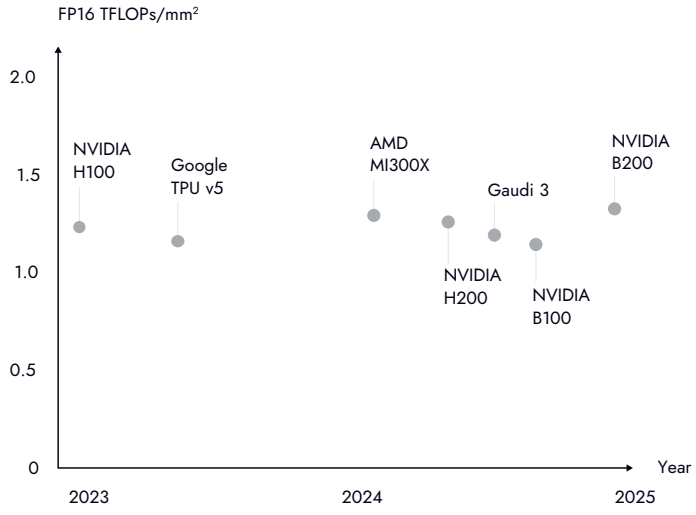

1 EE Times (2024); AI Needs New Breakthroughs in Energy-Efficient Computing 2 International Energy Agency (2025); Energy and AI Report
3 Semiconductor Insight (2024); AI Computing Hardware Market, Trends, Business Strategies 2025-2032 4 Etched, NVIDIA, Future Ventures

This rapid advancement is driven by three key levers:

- **Physical Scaling:** Photonic processors exploit massive parallelism through multiplexing across space, time and wavelength making them ideal for large-scale networks with high-throughput.
- **Arithmetical Scaling:** Photonics naturally supports high-throughput mathematical operations essential for AI, such as matrix-vector multiplications or Fourier transforms, which can be directly implemented in the optical domain.
- **Algorithmic Scaling:** By using algorithms tailored to photonic hardware, such as nonlinear functions, results equivalent to digital can be achieved with fewer computational cycles.

OPERATIONAL AT LEADING SUPERCOMPUTING CENTRES

The Leibniz Supercomputing Centre (LRZ) and the Juelich Supercomputing Centre (JSC) – two of Europe's leading HPC data centers – integrate Q.ANT's NPS into their operational HPC environment. These first deployments mark a major step towards redefining how data centers approach performance, footprint and energy-efficiency.

PHOTONIC COMPUTING VALIDATED IN PRODUCTION

Benchmarking of Q.ANT's first- and second-generation Native Processing Unit at the Leibniz Supercomputing Centre demonstrates significant gains in performance and efficiency under operational HPC conditions:

- **50× higher throughput for matrix multiplications**
- **25× faster inference** on a ResNet-18 convolutional neural network
- **6× lower energy consumption** for typical workloads
- **Enhanced analog processing units** optimized for nonlinear functions, reducing parameter counts and training depth
- **Accuracy** sufficient to support **state-of-the-art AI applications**.

PHOTONIC COMPUTING ENTERS THE COMMERCIAL CLOUD

With the deployment of Q.ANT's photonic processor at IONOS - a major European cloud and hosting provider - photonic computing moves from research HPC environments into commercial cloud infrastructure. As the first commercial customer, IONOS will make photonic AI acceleration accessible to enterprise users through its cloud ecosystem. This milestone demonstrates that photonic processors are a scalable path for data centers to address growing AI compute demands.

THE Q.ANT NATIVE PROCESSING SERVER NPS

The Q.ANT Native Processing Server NPS is a photonic analog processor as 19" rack-mountable server. Its core is the Native Processing Unit (NPU), a photonic AI Accelerator as PCIe card. Its Plug & Play system design enables a seamless integration into data centers and HPC environments, providing immediate access to photonic computing. The NPS is upgradable with additional NPUs to increase processing power as workloads grow.

BENEFITS FOR DATA CENTERS

Native Computing by Q.ANT promises:

- **Up to 30x higher energy efficiency** per application
- **Higher computational density** for increased throughput without expanding footprint
- **Reduced infrastructure complexity** and space requirements
- **Lower cooling requirements** through removal of on-chip heat
- **Reduced operational costs** through lower energy consumption and performance gains
- **Early access to next-generation compute architecture** enables pioneering research
- **Seamless deployment** as 19" server fully compatible with standard x86 software stacks

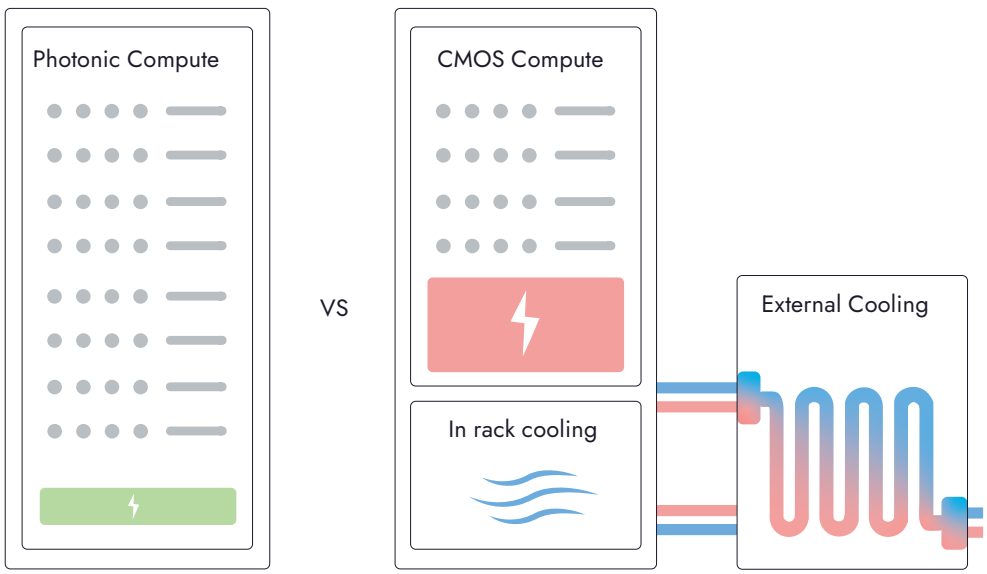

TECHNICAL DATA

DATA	VALUE/ DESIGNATION
Form Factor	19'' rack, 2U height Height: 87 mm; Width: 482 mm; Length: 889 mm
System Node	x86 processor architecture; supports multiple NPU cards per server
Operating System	Linux Debian/Ubuntu with long-term support
Network interface	2x 10 GbE RJ45, 1x 1 GbE RJ45 service interface, other optional
Software Interface	C / C++, Python API and PyTorch
API to NPU Subsystem	Linux device driver
Native Processing Unit NPU	Full length PCIe card with 2 slot height PCIe Gen4 x8 interface, shared memory & I/O windows Upgradable with enhanced photonic integrated circuits Upgradable with enhanced logic functions for performance & functional improvements
Throughput	8 GOPS
Photonic Integrated Circuit	Ultrafast photonic core based on z-cut Lithium Niobate on Insulator (LNoI)
Power Supply	Dual 2.000 W; AC input 115-240 VAC; 50-60 Hz; C13-to-C14/CEE7-VII or country specific
Power Consumption NPU	150 W
Operating Temperature Range	15 to 30°C
Weight	NPS (without NPU cards): 38.8 kg ; NPU: 2.38 kg

RECOGNIZED BY GARTNER

> Photonic computing has several potential benefits over electronic computing, including increased bandwidth, processing power and storage, all while keeping energy and power consumption under control.

Gartner Research (2025): Emerging Tech: Emergence Cycle for Generative AI

Q.ANT is recognized as a Sample Vendor for Photonic Computing in the Gartner® Hype Cycle™ for Data Center Infrastructure Technologies report 2025.

GARTNER is a registered trademark and service mark of Gartner and Hype Cycle is a registered trademark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and are used herein with permission. All rights reserved. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

PHOTONIC INTEGRATED CIRCUIT AT THE HEART OF NPS

Q.ANT's proprietary Photonic Integrated Circuit PIC based on Thin Film Lithium Niobate is at the heart of this innovation. It offers precise light control at the chip level. Q.ANT controls the entire value chain from raw materials to fully functional systems to make these processors achieve superior mathematical and algorithmic performance.

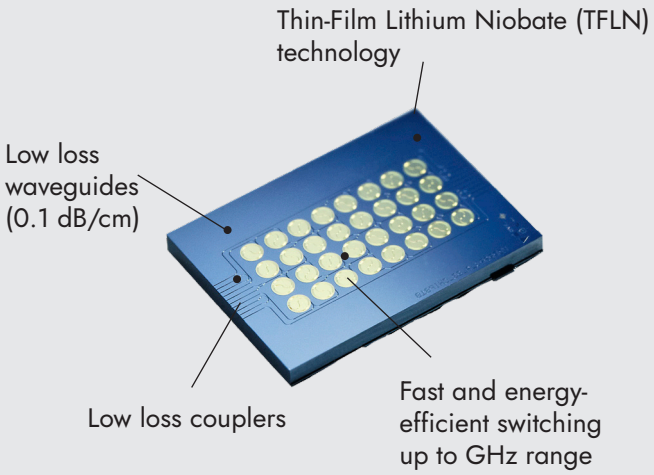

CORE MATHEMATICAL OPERATIONS ARE NATIVE IN OPTICS

Mathematical operations required for a wide range of AI can be implemented directly in the optical domain using dedicated elements on the photonic chip. Photonic computing requires fewer elements than CMOS-based architectures for the same basic mathematical functions, leading to reduced circuit complexity, lower latency and higher computational density.

Mathematical operation	**Required elements in optics**	**Required number of transistors in CMOS for 8-bit integer precision**
Summation	Q.ANT photonic element γ	Approx. 200 transistors
Subtraction	Q.ANT photonic element γ	Approx. 200 transistors
Multiplication	Q.ANT photonic element Ω	Approx. 1200 transistors
Root extraction	Q.ANT photonic element γ	Approx. 7700 transistors
Squaring	Q.ANT photonic element Ω	Approx. 1200 transistors
Fourier transform	Q.ANT photonic element X	Algorithm dependent, several millions for near real-time Fast Fourier Transforms

FROM WAFER TO SERVER – Q.ANT CONTROLS THE FULL VALUE CHAIN

Q.ANT manages the entire value chain in photonic computing, ensuring full control over performance, quality, and scalability. The photonic chips at the heart of NPS are produced at Q.ANT's own pilot line specifically designed for chip production using Thin Film Lithium Niobate (TFLN), the optimal material for photonic computing. By upcycling an existing CMOS production line at IMS CHIPS, this partnership establishes a blueprint for democratizing chip production capacity ensuring manufacturing resilience and reducing dependency on global supply chains.

IMS CHIPS

TRANSISTORS DEFINED THE PAST. LIGHT DEFINES THE FUTURE.

One Year in Photonic Processing Development equals Ten Years in CMOS Development:

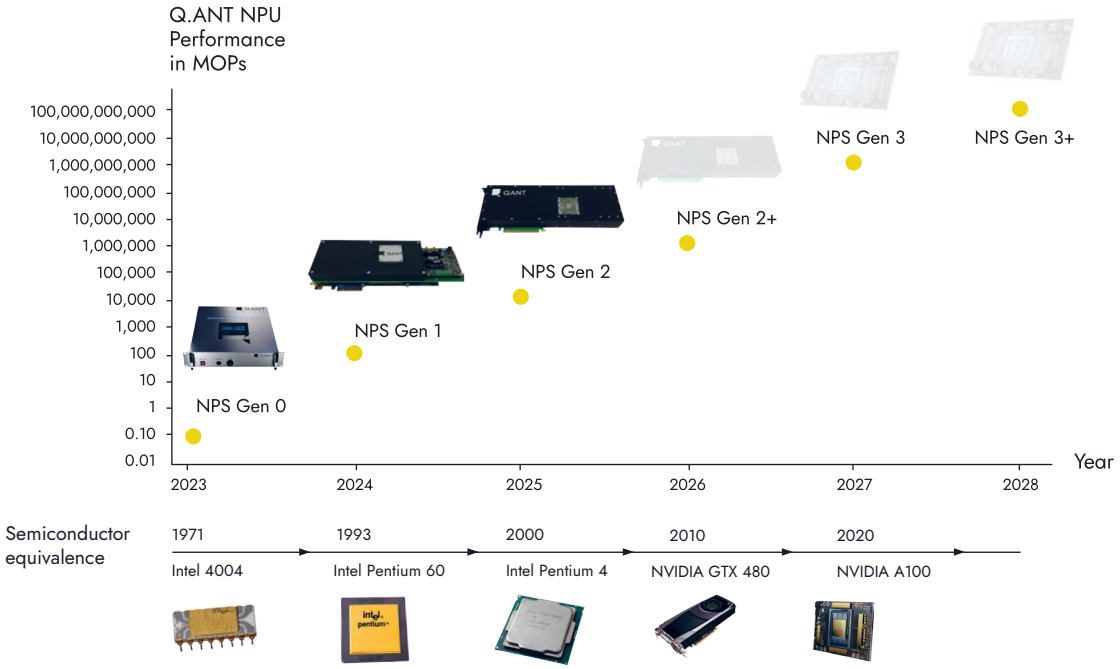

THIN FILM LITHIUM NIOBATE – THE MATERIAL OF CHOICE FOR PHOTONIC CHIPS

Q.ANT builds its photonic processors on a proprietary Thin-Film Lithium Niobate (TFLN) platform. This material system enables the photonic integrated circuits (PICs) at the core of our NPUs. On a silicon wafer, a thin layer of lithium niobate is bonded to create optical waveguides, modulators, and other functional blocks — allowing high-speed, precise control of light on a single chip.
PICs based on TFLN provide the physical foundation for scalable, efficient photonic computing:

- No thermal crosstalk between optical components since electrical fields, not heat, change the material properties.
- Ultra-Fast modulation and switching speeds of few tens of GHz.
- Native non-linearities offer more complex functions for computing.
- Room temperature operation opens up a wide range of application fields.

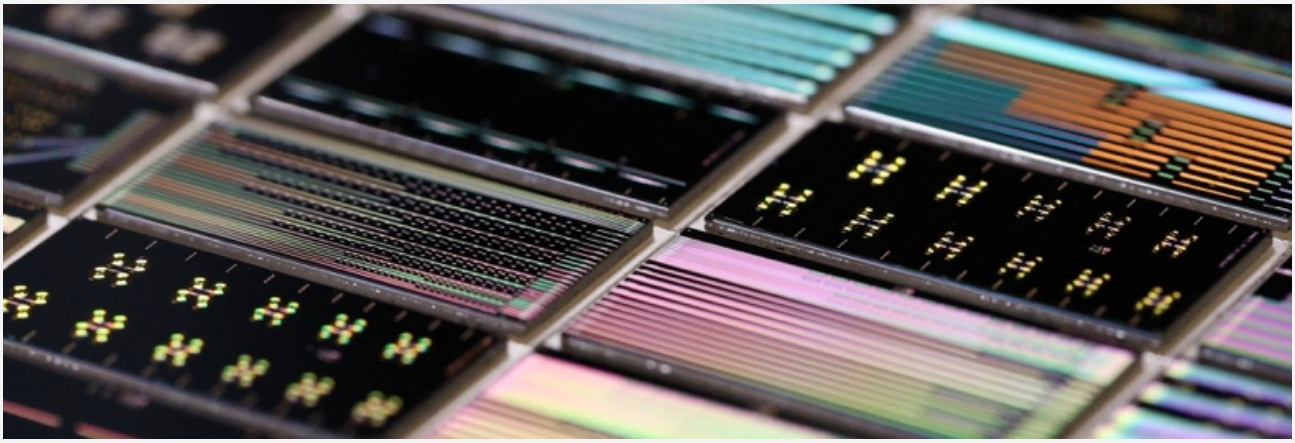

Q.ANT NPS – THE CHOICE FOR EFFICIENT NONLINEAR NETWORKS

While CMOS processors excel at linear, sequential processing, photonic processors are the natural hardware fit for large-scale nonlinear algorithms:

- **Fewer parameters:** Networks with nonlinear functions reduce the number of model parameters needed allowing higher accuracy per parameter and training budget.
- **Fewer operations:** In photonic processors one single optical element performs one nonlinear operation while CMOS requires 100-1000 transistors and multiple cycles.

The Q.ANT Native Processing Unit NPU is the Choice for efficient nonlinear Networks

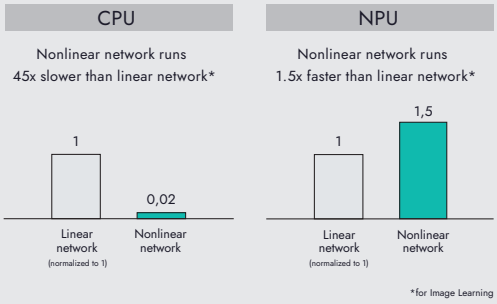

PHOTONIC COMPUTING WILL BENEFIT A BROAD APPLICATION SPACE

Photonic Computing has the potential to unlock new levels of performance and energy efficiency for a variety of applications through the efficient use of nonlinear functions.

Artificial Intelligence	Function Approximation	Nonlinear Networks for high-frequency Signals	Nonlinear functions are native in optics, reducing the number of model paramenters	Financial Modeling Sales Forecasting National Grid Forecast
	Natural Language Processing	Transformer based Language Models	More efficient model architecture with fewer parameters for accelerated training & inference	Agentic AI & Chatbots Text Generation Programming
	Image Processing	Convolutional Neural Networks	Faster computations by transforming convolutions to multiplications in Fourier domain, reducing energy and memory usage	Computer Vision Medical Image Analysis Quality Control in Manufacturing
Computational Modeling & Simulation	AI-enhanced PDE Simulation	Neural Operators and Deep Operator Networks	Faster and more energy-efficient simulations remaining robust to deviations from the ideal Fourier transform	Finite-Element Analysis Computational Fluid Dynamics Process Modeling
	Molecular Property Prediction	Graph Neural Networks	Faster and more efficient networks with fewer parameters while maintaining high accuracy	Drug Discovery Materials Design Bio and Chem Informatics

WHITEPAPER: EXPLORING EARLY AI APPLICATIONS ON Q.ANT'S PHOTONIC NATIVE PROCESSING SERVER

In this recent whitepaper, we explore early AI applications running on Q.ANT's photonic Native Processing Units — from neural network-based classification and segmentation tasks using pre-trained weights, to image learning with advanced, non-standard neural network architectures. These experiments demonstrate the real-world capabilities of photonic computing today.

Access the Whitepaper:

CLIMBING THE COMPLEXITY LADDER OF AI MODELS

Each generation of Q.ANT's photonic processor demonstrates a consistent increase in AI model complexity and processing performance - providing clear evidence that photonic computing can reliably execute basic AI workloads and scale toward future generations of AI.

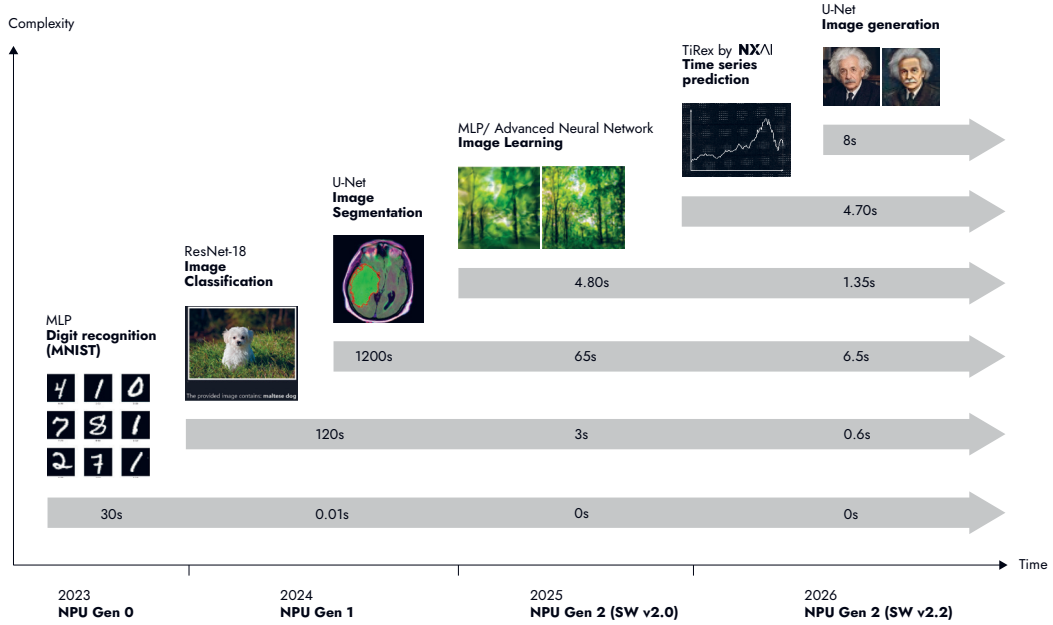

THE CLEAR PICTURE – IMAGE LEARNING ON Q.ANT NPU

In this example, a network using learnable nonlinear functions on Q.ANT's NPU reconstructs complex image patterns more accuretaly than a linear network on a CPU while 2x less parameters and 3x less operations are needed.

Original image

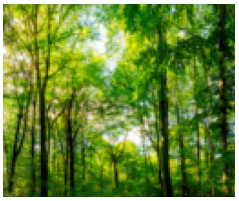

Linear network on CPU

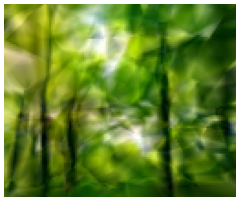

parameters: ~20k
operations: ~670m

Nonlinear network on NPU

parameters: ~10k
operations: ~250m

PROVEN PRECISION FOR STATE-OF-THE-ART GENERATIVE AI

A key challenge for photonic analog computing has been the question of precision in large neural networks. A 59-million-parameter diffusion model for image generation successfully runs on a Q.ANT Native Processing Unit - proving that error propagation is not a limiting factor for photonic AI computing. Support for industry-standard bfloat16 precision allows existing models to run without adaptation, validating the technology for real-world AI workloads.

Original

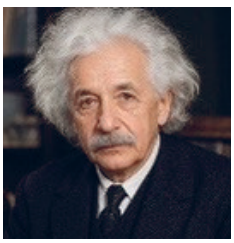

Generated image in pixel art

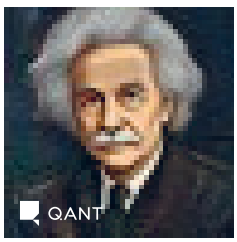

Q.ANT INTEGRATES INTO THE ESTABLISHED COMPUTE LANDSCAPE

Q.ANT's photonic processing solution seamlessly integrates into the existing compute landscape. The Native Processing Unit at the heart of the NPS provides a PCIe interface housed in a standard 19" server, which makes the system plug-and-play. The NPU can be accessed via a software interface with C/C++ and Python APIs and will integrate into common AI frameworks such as PyTorch. Q.ANT supports customers in creating custom applications, providing the Q.ANT Photonic Algorithms Library and training resources.

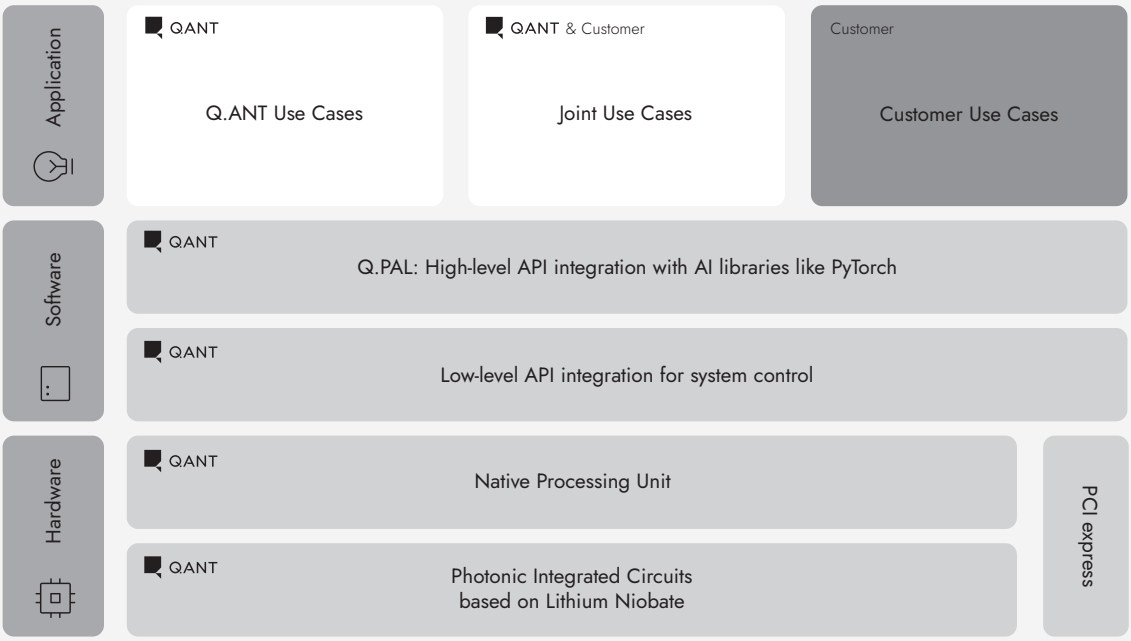

A DAY IN THE LIFE OF AN AI MODEL

Based on a standard neural network

Applications

Application codes in PyTorch:

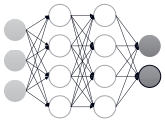

Standard libs and frameworks

PyTorch generates input to Compiler: $Z = AW + b \; ; \; A = \sigma(Z)$

Driver and base interface

Toolkit submits through driver for optical implementation
e.g. multiply elements of:

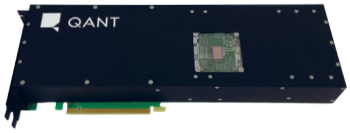

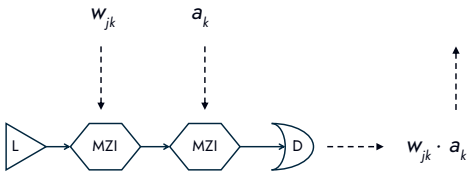

Note: MLIR code is not functional and just for illustration purposes

Internal | © Q.ANT GmbH 2026

Q.PAL – THE Q.ANT PHOTONIC ALGORITHMS LIBRARY

The Q.ANT Photonic Algorithms Library Q.PAL is the software interface to the NPU which enables users to operate directly at the multiplication level or to leverage optimized neural network operations such as fully connected layers or convolutional layers. Q.PAL offers a comprehensive collection of example applications that illustrate how AI applications can be enhanced. These examples can be implemented directly or as a foundation for creating custom use cases.

Name	Description	Programming Language
Matrix Multiplication	Classification of an image (e.g. based on ImageNet data set)	Python / C++
Image Classification	Classification of an image (e.g. based on ImageNet data set)	Python (Jupyter)
Semantic Segmentation	Segmentation of an image (based on a brain MRI data set)	Python (Jupyter)
Complex Line Fitting	Fitting of a high frequency line with a nonlinear network (e.g. based on simulated training data)	Python (Jupyter)

Q.PAL FUNCTIONALITY

The Q.ANT Photonic Algorithms Library provides developers with a powerful operation library that encompasses both native hardware operations and complex, higher-level functions built upon them — enabling seamless integration into modern AI workloads.

Native operations

- Element-wise vector multiplication
- Matrix-vector multiplication
- Matrix-matrix multiplication
- Cosine on vector and element-wise multiplication with a weight vector
- More native operations to come...

Non-native operations (building on native operations)

- Forward propagation layer of a neural network (A*x multiplication, +b sum)
- Convolution layer forward pass
- ReLu/Sigmoid/Softmax/different pooling layer forward passes

All operations are fully supported in bfloat16, enabling efficient computation with reduced memory footprint while maintaining the numerical stability required for AI workloads.

Q.ANT

JOIN THE NEXT ERA OF COMPUTING WITH LIGHT!

CONTACT

native-computing@qant.com
+49 711 25245-0
www.qant.com

Q.ANT GmbH
Handwerkstraße 29
70565 Stuttgart, Germany