Press Release

Q.ANT Runs Generative AI on Photonic Hardware

At ISC High Performance 2026: Generative AI and Recurrent Networks run on Q.ANT's Second-Generation Photonic Processor

Stuttgart – June 23, 2026 – Q.ANT, the pioneer in commercial photonic computing, today demonstrated the first complex, production-relevant AI workloads on its photonic hardware.

Q.ANT successfully demonstrated a diffusion model and a recurrent neural network on its second-generation Native Processing Unit (NPU) at ISC High Performance 2026 in Hamburg. This proves that Q.ANT's photonic architecture supports the full breadth of modern AI capabilities in generative image synthesis and sequential time series prediction.

The ISC demonstrations build on ecosystem progress already underway. Earlier this year, independent developers at Daisytuner compiled and deployed an object detection model directly from PyTorch onto Q.ANT's photonic processor. This marked the first time an AI model from a standard ML framework has been successfully compiled for photonic hardware.

These developments reveal that Q.ANT's NPS has advanced beyond foundational algorithms to genuine commercial applications. Using Q.ANT hardware, these high-performance computing tasks target to operate with 30x the energy efficiency of classical processors in equivalent matrix operations at the photonic circuit level.

"Q.ANT's photonic architecture changes the energy calculus for AI infrastructure." says Q.ANT founder and CEO, Dr. Michael Förtsch. "When you perform computation with light instead of transistors, you reduce energy consumption at the source. And every serious conversation about the future of AI acknowledges that energy is the bottleneck the industry must break through. Our recent demonstrations of generative AI show that photonic hardware can carry the mathematical load of the most demanding modern AI workloads."

Generative AI on Photonic Hardware

To illustrate its capabilities for generative AI, Q.ANT hardware ran a diffusion model to perform image-to-image synthesis – a class of generative AI workloads defined by iterative, parallelized matrix operations. This image-to-image application highlights the viability of photonics in one of the most computationally intensive neural network architectures in modern AI and marks the first time a diffusion model of this complexity has run on photonic hardware.

Diffusion models generate images through repeated forward passes of a deep neural network in dense matrix operations. By executing the primary computational layer using light instead of transistors, Q.ANT's photonic processor moves beyond foundational algorithms into processing the linear arithmetic at the core of modern AI applications.

"Diffusion models are widely used and computationally demanding approaches in modern generative AI. They rely on repeated, large-scale operations to gradually produce a coherent output," says Professor Dr. Björn Ommer, head of the Computer Vision & Learning Group at Ludwig Maximilian

University of Munich (LMU) and the leading researcher behind the development of the stable diffusion model. “If photonic hardware could execute such workloads efficiently and reliably, it would be an exciting indication that alternative computing substrates may play an important role in the future of generative AI.”

Time Series Prediction with xLSTM

Proving the extensibility of its hardware in structurally distinct architectures, Q.ANT also executed the TiRex time series prediction model developed by NXAI, the Austrian frontier AI lab commercializing the Extended Long Short-Term Memory (xLSTM) architecture for enterprise applications.

"We've been pushing the balance of performance and power consumption with TiRex from the beginning. Seeing it run on Q.ANT's photonic hardware is amazing and opens a new chapter," says Lukas Fischer, Head of Applied Research at NXAI. "xLSTM architecture on photonic systems could redefine what energy-efficient AI even means."

Unlike transformer-based models, xLSTM is a recurrent neural network designed to identify patterns across sequential data and predict future values over long-time horizons. NXAI’s commercial TiRex model, with production-tuned weights, targets financial market analysis, supply chain optimization, weather forecasting and traffic flow simulation.

Through the xLSTM and diffusion model demonstrations, Q.ANT now shows that its hardware can operate on the most demanding classes of modern AI. Running both shows that Q.ANT hardware is built for the breadth of AI use cases.

Building a Photonic Computing Ecosystem

The ISC demonstrations are the latest in a series of third-party integrations, commercial partnerships and institutional deployments for Q.ANT.

In May, Q.ANT announced the first commercial orders for its hardware through a partnership with the German cloud services provider, IONOS. And in April, Q.ANT partner Daisytuner, revealed the development of a compiler using standard AI toolchains for a live object detection application. Leading institutions like the Leibniz Supercomputing Centre Munich and Jülich Supercomputing Centre, two of Europe's breakthrough HPC facilities, are running Q.ANT hardware in live production.

This ecosystem forming around Q.ANT's platform reflects the institutional confidence in photonic computing as the infrastructure for next generation compute.

About Q.ANT
Q.ANT is commercializing photonic accelerators for AI and high-performance computing, offering a scalable alternative to transistor-based systems. Its Native Processing Units (NPUs) use photonic integrated circuits based on a lithium niobate material platform to perform mathematical operations directly on the chip using optical signals, enabling energy-efficient co-processing for complex

computational tasks. Q.ANT operates its own TFLN chip pilot line in collaboration with IMS CHIPS. Q.ANT was founded in 2018, and is headquartered in Stuttgart, Germany.

++++++++

Pictures

IMG 1

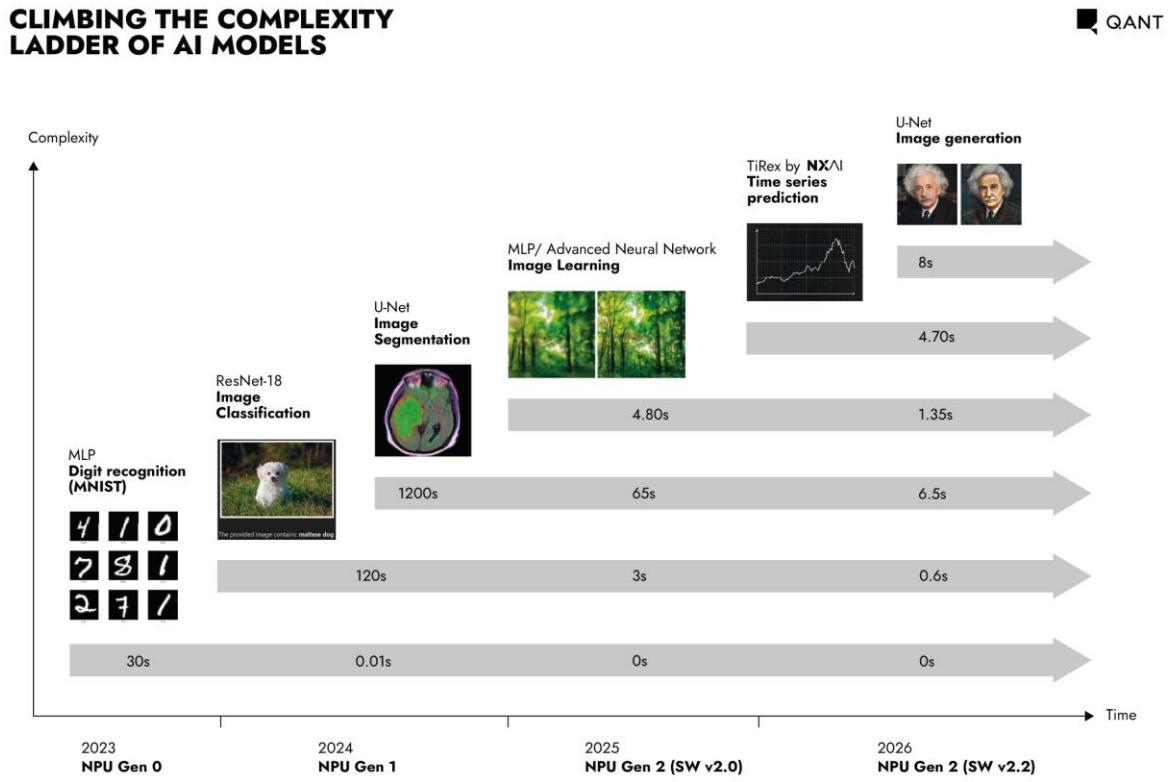

IMG 2

IMG 3

Captions

1. **Climbing the complexity ladder of AI Models**
2. **Q.ANT Native Processing Unit NPU 2**
3. **Q.ANT Working on the Diffusion Model**

Media Contact
LaunchSquad for Q.ANT
qant@launchsquad.com

Q.ANT
Jörg Kochendörfer
joerg.kochendoerfer@qant.gmbh
+49 160 5619730