



with the Aurora2 OPU

LightOn Appliance, a photonic co-processor sets a new pathway for Transformative AI computing

LightOn Appliance is the world's most powerful photonic co-processor for AI and HPC. Initially available on its own cloud platform, LightOn's technology can now be installed on-premises, as a 2U rackable unit. LightOn's Optical Processing Unit (OPU) technology combines the **massively parallel scales** required by today's most challenging AI workloads, with an **extremely low power consumption**. LightOn Appliance is today's most accessible way of discovering the future of computing, in your own datacenter.

Peta-scale performance, ultra-low power consumption

The LightOn Appliance leverages photonics to **process large-scale generic data into information-revealing "sketches"** opening new avenues ranging from fundamental issues in mathematics and machine learning, up to alleviating some of today's largest bottlenecks in production for data science, **HPC** and **AI**. Current specifications are setting new standards for sustainable AI acceleration, with more than **1.5 PetaOPS** at **30 W TDP**.

Offloading some memory-intensive dense computations on the Aurora2 OPU makes your CPU and GPU more efficient, and results in significantly **accelerated computing pipelines**, sometimes by a factor of 10 or more.

The OPU is **seamlessly integrated** into your Python / PyTorch workloads: one line to import our Python API, and a single command to process a block of data on the OPU. All this compliant with the most popular Python-based ML packages such as PyTorch and scikit-learn.

LightOn Appliance supports the following tools:



System Specifications

OPU type	Aurora2
Photonic core	One Nitro gen2 photonic core
Operation	Dense random layer
Max Input size	1 million (1 039 680) up to 6-bit
Max Output size	2 million (2 267 392) 8-bit
Operating modes	Batch or Online
Peak batch mode performance (non-linear)	50 TOPS/W* (up to 1 900 dense matrix vector multiplications per second)
Peak batch mode performance (linear)	10 TOPS/W* (up to 360 dense matrix vector multiplications per second)
Frameworks compatibility	PyTorch 1.0+, scikit-learn
SDK / Library	Appliance SDK/LightOnML v1.4
Peak TDP	30 W
Operating environment	Actively cooled data center
Support	Remote support on-site support on demand
	Installation support available on demand (HW and SW)
Minimum requirements	Intel Xeon x86_64 server-grade CPU***
	Ubuntu 18, Debian buster, RHEL 8
	Python 3.7+
	One PCIe Gen2 x4 single-width slot for a low profile card
Form Factor	2U
Power supply	125/250V 50/60Hz


* Future updates will improve performance

** Not tested on AMD CPUs but it is expected to work as is

Unparalleled performance across a wide variety of applications

LightOn Appliance for AI acceleration

Approximating Kernel Ridge Regression for classification tasks



Performance gains on Kernel ridge approximation for classification tasks. Dataset qm7 (quantum chemistry), high energy physics, and others. The OPU is compared to an NVIDIA P100 GPU (250 W). GPU RAM limit was hit at 32GB. Results acquired extrapolating to 1M features. OPU: Aurora 1.5 (30 W).

Real-time AI analysis of large-scale HPC results

Change detection in Molecular Dynamics



15x faster than FastFood on CPU at 50k atoms. For 700k + atoms, NEWMA RP on OPU is expected to be 30x faster than NEWMA FF on CPU. Library: LightOnML, Dataset: Molecular Dynamics simulations (HPC, Anton), OPU: Aurora 1.5

LightOn Appliance has been successfully used with a variety of ML algorithms, including:

- Transfer Learning
- Kernel Ridge Regression
- Sketching Algorithms
- Randomized Linear Algebra
- Reinforcement learning
- Locality sensitive hashing

... and with the following data types:

- images and videos
- time-series
- text / tokenized text
- graphs
- sound signals
- HPC simulation results
- scientific experimental data

For more use-cases and benchmarks visit:

- lighton.ai/lighton-white-paper
- github.com/lightonai/opu-benchmarks

A tool for the community

LightOn Appliance is LightOn's first on-premises OPU product for ML and HPC. Leasing package includes access to an Aurora2 OPU built upon the Nitro2 photonic core.

For more information, leasing options, and specific enquiries visit lighton.ai/lighton-appliance/ or contact us at client.relations@lighton.ai

A community of users. We organise monthly meetups with world-class guest researchers. We open-source most of LightOn AI Research (LAIR) algorithms and present our research at major international conferences. Please check our publications page: lighton.ai/publications

LightOn's Technology has been featured at:



IEEE HotChips '20 / '21



AI Hardware Summit '20



NeurIPS '19 / '20 / '21



IBM / IEEE AI Compute Symposium '20



SuperComputing '19 Optical Computing workshop




IEEE International Conference on Rebooting Computing ICRC '19



HUAWEI Technology Innovation Forum '19 & '20



ATOS Annual Expert Community Convention '18



MATLAB EXPO '19