

Blog Posts and Papers Published During the First Months of 2026

Robotics AI Suite Races Forward on Panther Lake and Bartlett Lake	1
Power Up Patient Monitoring with the Health and Life Sciences AI Suite on an Intel Platform	1
Open Edge Platform Optimized for Running Edge AI on Panther Lake	1
Other Blog Posts and Papers on Edge AI During the First Months of 2026	1
Find Open Edge Platform on GitHub	3
Find My Work on LinkedIn and Medium	3

Robotics AI Suite Races Forward on Panther Lake and Bartlett Lake

This comprehensive [Medium.com](#) and [GitHub blog post](#) highlights the Robotics AI Suite, its capabilities, and its sample applications. As an overview, the post includes sections on the components of the 2026.0 release, performance testing, partnering, Panther Lake and Bartlett Lake, and links to key Intel resources for robotics.

Power Up Patient Monitoring with the Health and Life Sciences AI Suite on an Intel Platform

The release of the Health and Life Sciences AI Suite and its Multi-Modal Patient Monitoring app demonstrates how to run multiple AI workloads concurrently on a single Intel-powered edge device without a discrete GPU so that you, as a medical AI developer, can see the cost-effective performance of Intel® Core™ Ultra processors in action. This [GitHub blog post](#) explains how to power up patient monitoring at scale.

Open Edge Platform Optimized for Running Edge AI on Panther Lake

Open Edge Platform and nearly all its infrastructure, libraries, suites, microservices, and applications have been tested and tuned as fully functional for performance on the Panther Lake architecture. With Open Edge Platform now optimized for Intel® Core™ Ultra series 3 processors, code-named Panther Lake, the Open Edge Platform AI suites, libraries, tools, and frameworks deliver a unified silicon, platform, and software stack with cost-effective performance for AI workloads at scale. [This matters; here's why.](#)

Other Blog Posts and Papers on Edge AI During the First Months of 2026

Here's a quick summary of some of the other blog posts and papers I've published so far in 2026, typically as part of highlighting the capabilities, use cases, and business benefits of [Open Edge Platform](#) from Intel or the cost-effective performance of Intel processors for AI workloads.

- [Light Up the Multi-Modal Future of Edge AI with Open Edge Platform 2026.0](#): As edge workloads rapidly evolve to include AI models and capabilities for sensing conditions, making sense of those conditions, and acting in some way on that data, the applications – and their underlying processors – must remain predictable and performant at scale.
- [Build and Evaluate Robotics AI Workloads on the AAEON CEXD-INTRBL Reference Board](#). As a robotics developer, you typically need to stitch together cameras, ROS 2, AI acceleration, sensor interfaces, and deployment tools before you can validate an application. The CEXD-INTRBL robotics reference board from AAEON gives you a pre-integrated path to prototype, evaluate, and tune robotics AI workloads on Intel® Core™ Ultra Series 3 processors, code-named Panther Lake. The reference board combines an Intel® Core™ Ultra X7 Processor 358H, integrated Intel® Arc™ GPU, and Intel® NPU 5.0 in an embedded development system designed for AI-powered robotics workloads that require real-time perception, decision-making, and control at the edge.
- [Go Edge Native in Style with the Edge AI Handbook and Panther Lake](#). Benchmark results that cover AI tasks performed at the edge help you make decisions about initial system selection and setup, including for both bare-metal and containerization methods.
- [Edge AI Handbook: A Practical Approach to AI Application Development](#). The Edge AI Handbook: A Practical Approach illustrates how to develop AI applications. The handbook spans system selection, system setup, system profiling, accuracy optimization, performance optimization, and deployment.

Sample applications serve as implementation references, and you can use the handbook to set up and optimize such systems as Deep Learning Streamer, Geti™, and OpenVINO™. Benchmark results that cover AI tasks performed at the edge help you make decisions about initial system selection and setup.

- [Fast Track Media Analytics with Videos on Deep Learning Streamer](#). Deep Learning Streamer is an open-source streaming media analytics framework. Based on the GStreamer multimedia framework, DL Streamer lets you create media analytics pipelines for the cloud or at the edge.
- [Build Better Vision Models Faster with Geti Video Demos on YouTube](#). Check out the latest video demonstrations of Geti™ software to find out how to build computer vision models in a fraction of the time and with less data, in effect turbocharging the entire vision AI lifecycle, including training on Intel® Arc™ GPUs. Geti software eases laborious data upload, labeling, model training, retraining, and optimization tasks to streamline the development of computer vision models. You can rapidly build vision models for various processes, including detecting defective parts in a production line, reducing downtime on the factory floor, or automating inventory management.
- [Video Demos of the Smart Parking and Smart Intersection Apps from Metro AI Suites](#). The Smart Intersection video on YouTube uses OpenVINO for inference, Deep Learning Streamer for full video analytics, and SceneScape for 3D live scene modeling. The Smart Traffic Intersection sample application, also part of the Metro AI Suite, combines analytics from multiple traffic cameras to provide a unified intersection view, enabling advanced use cases such as object tracking across multiple viewpoints, motion vector analysis (e.g., speed and heading), and understanding object interactions in three-dimensional space. To find out more, check out the recent Medium blog post titled [Intersections to Infrastructure: Enabling Edge AI with Metro AI Suite](#).
- [Edge AI Suites Optimized and Tested for Running Edge AI Applications on Panther Lake](#). The Robotics AI Suite enables you to take your own advanced AI capabilities and use them to rapidly develop cost-effective robots using Panther Lake for both controls and AI-based perception. You can use the various Edge AI Suites to benchmark the performance of your own edge AI workloads on Intel® Core™ Ultra series 3 processors.
- [Go Bigger Faster: Performance Advances and KPIs Tied to Release 2025.2](#). The latest release of Edge Manageability Framework demonstrates performance advances for edge AI workloads by publishing KPIs, including expanded support for performance assessments of reference applications like pallet defect detection, image-based video search, chat Q&A, and smart parking. These performance indicators help you, as system builders and DevOps engineers, assess deployment scenarios on Intel processors to ensure that you can take advantage of the latest hardware capabilities.
- [Did you know you can quickly deploy AI models with Edge AI Demo Studio on edge-optimized Intel hardware?](#) Edge AI Demo Studio lets you rapidly deploy and manage AI models on Intel-optimized hardware primed for edge computing. The demo studio is a modern toolkit for deploying, managing, and serving AI models from Hugging Face on edge platforms. The demo studio lets you use a web UI to manage workloads, making it easy to download, convert, and serve Hugging Face models with minimal setup. You can use various AI services to power your applications, including text generation with large language models (LLMs), text to speech (TTS), speech to text (STT), image generation, MCP Manager, and wake word detection. A sample use case is RAG chat.
- [Did you know you can transform your Intel platform into a powerful edge AI solution with ready-to-deploy use cases?](#) Here's how to streamline the deployment of an edge AI use case on an Intel platform.
- [Did you know you can set up your Intel hardware to validate Edge AI use cases in about 5 minutes?](#) Here's the fast path for setting up Intel hardware to validate a variety of Edge AI use cases with Edge AI Suites and, optionally, other components of Open Edge Platform.
- A Medium blog post titled [Validating the Performance of Edge AI Workloads on Intel Processors with Edge Microvisor Toolkit](#) on running easy-to-benchmark sample applications for smart traffic intersections and wind turbine predictive maintenance on Edge Microvisor Toolkit to showcase the

power and performance of running edge AI workloads on Intel processors.

- [Overview of Open Edge Platform 2025.2: Suites, Libraries, Toolkits, and Frameworks to Propel Edge AI into 2026](#). Open Edge Platform from Intel delivers a modular, composable stack of open-source software optimized to rapidly deploy, secure, manage, and scale AI solutions at distributed edge sites. To blaze an edge AI trail for the new year, we released Open Edge Platform 2025.2 in December. The release lights up the first few months of 2026 with AI models, sample applications, SDKs, microservices, edge infrastructure, performance-optimized pipelines and workflows, AI-optimized code, benchmarking tools, and industry-specific libraries. The newly minted Robotics AI Suite takes flight, and the latest Geti release lets you fine tune vision models on Intel® Core™ Ultra Series 2 processors with a built-in GPU so you can find your own way with ease.
- [Open Edge Platform Optimized for Running Edge AI on Panther Lake and Other Edge AI Blog Posts](#). This [LinkedIn article](#) is a quick summary of the blog posts and papers I've published so far in 2026, typically as part of highlighting the capabilities, use cases, and business benefits of Open Edge Platform from Intel, its AI suites, or the cost-effective performance of Intel processors for AI workloads.

Find Open Edge Platform on GitHub

A modular, composable stack of open-source software optimized to rapidly deploy, secure, and scale AI solutions at distributed edge sites, [Open Edge Platform](#) comprises [Edge Microvisor Toolkit](#), [Edge Manageability Framework](#), [Edge AI Libraries](#), and [Edge AI Suites](#). See, for instance, my Medium blog post on [Evaluating Security Capabilities for Edge AI](#).

In addition, the [OS Image Composer](#) tool lets you rapidly create and easily manage custom Linux images to meet your organizations' exacting requirements for edge AI workloads. The tool gives you a standardized, extensible framework for composing operating system images from pre-built artifacts of any Linux distribution that supports Debian or RPM packages. See my recent blog post on Medium: [Composing Custom Linux Images with the Inaugural Release of OS Image Composer](#).

Find My Work on LinkedIn and Medium

I'm a writer, editor, and AI educator working at the interdisciplinary nexus of critical thought, text, and technology. I have a master's degree in Linguistics from The Graduate Center, City University of New York, and have taught classes at Bronx Community College, Harlem Hospital, Maimonides Medical Center, the Brooklyn Public Library, and the New York City College of Technology (City Tech). More recently, I have worked as a technology writer, editor, or educator at AT&T, Verizon, Reed Elsevier (now RELX), Microsoft, EMC, VMware, Broadcom, and Intel. I write about technology, linguistics, AI/ML, corpus analysis, and natural language processing, or NLP.

- [Medium.com Page](#).
- [LinkedIn Page](#).
- [YouTube Page](#).
- [Profile on Google Scholar](#). (My essay on Max Weber was cited in the 2002 APSA Presidential Address: The Public Role of Political Science by Robert D. Putnam; you can also find Putnam's address on Google Scholar. In 2003, Encyclopaedia Britannica named my essay [best of the web for social science](#).)
- [Resume](#).
- [Publishing Portfolio: Technology Writing and Technical Marketing](#).
- [Blog Posts and Papers on Edge AI for 2025](#).
- [List of Publications on Technology](#).
- [More Edge AI blog posts on Medium by Steven Hoenisch](#).