

HPEC + AI = Predictive Maintenance ROI

When you peel away all the sophistication and innovation of technologies like predictive maintenance, they are, at their core, cost-saving solutions.

Peter Darveau is Head of Engineering and Principal Engineer at [Hexagon Technology Inc., a technical services firm that specializes in automation systems, safety-instrumented devices, and machine learning.](#)

He acknowledges that the first goal of predictive maintenance is to analyze the behavioral patterns of a machine over time to prevent system failures and optimize performance. The ultimate goal, of course, is more agile, competitive, and profitable operations.

“If you’ve got a \$5 million piece of factory equipment, you finance 80% of it at a 5% interest rate, which is pretty typical, and you just extend its life for one year, you’re saving \$200,000,” Darveau says. “If you can just get one more year, it’s an attractive payback. So right away, there’s interest.”

But interest doesn’t guarantee simplicity.

Feature Engineering: A Path to Operational AI

Automating the process of predicting equipment failures and recommending actions to prevent them requires AI. And for AI to be effective, humans must lend a hand.

Engineers have to generate data sets comprising equipment’s “normal” and “fault” conditions for AI algorithms to identify anomalies. What’s normal and what’s faulty are extracted from analog signals like vibration or acoustics and then classified accordingly through a process called feature engineering. But it’s a manual procedure

that requires extensive signal processing expertise to appropriately extract, evaluate, and classify signal data from the target machine.

The complexity of feature engineering explains why it’s taken so long for predictive maintenance strategies to be deployed.

Hexagon Technology helps companies implement asset monitoring solutions for large-scale industrial equipment like steel milling machines that handle multi-ton fragments of raw materials. Whenever a mass of materials moves, it vibrates the machine’s foundations, which were designed to accommodate enormous weight. Hexagon develops what’s called prognostics and availability monitoring (P&AM) systems that continuously analyze these structures so that integrity is maintained.

“In the past when we’ve done vibration analysis, we had to do so much just to have a useful data set,” Darveau explains. “If you wanted to analyze a noisy signal, you had to go through tons and tons of calculations. You had to do Fast Fourier Transforms (FFTs)—a method for transforming a function of time into a function of frequency—to filter out the noise, and then then you had to condition the data.”

Today, Hexagon can bypass much of the feature engineering process thanks to advances in high-performance embedded computing (HPEC). Specifically, the company is using Intel® Xeon® processors, Intel® Movidius™ VPUs, and an inference engine built on the Intel® OpenVINO™ Toolkit to find patterns in streaming waveforms instead of extracting features from complex analog signals.

“Because of the higher performance now on Intel® processors, we can do inferencing of video,” Darveau says. “The fact that it’s a video and I can chop up my frame means that instead of taking raw data and looking at very specific areas, I can easily look at patterns that happen over time and match it up with data from a simulation. Then we let the inference engine do its work.”

The Living Design of AI and ML Models

Hexagon’s P&AM solution is technically a functioning prototype that has been operationally deployed for two years at a plant like the one described above. During that time, Hexagon has more than doubled its video analytics performance from 30 to 60 FPS (or what the human eye can see) to 140 FPS. The system’s ability to accurately detect observed features has also increased to 95%.

Still, Darveau is aware that his system must achieve accuracies on the order of 99.999% before it’s ready for widespread commercial deployment.

“Because this environment is dynamic, and this is true for any AI or machine learning-type model, the model is a living thing. It is for us to make improvements to the system,” he explains. “And it’s going to be the last mile that’s going to be the toughest. We anticipate updating this model several times until we feel that it’s efficient enough.”

Knowing that their P&AM solution would be a living design from the start made the selection of OpenVINO an easy one.

Not only do its optimization features maximize execution performance on today’s processors, but its portability across CPUs, GPUs, FPGAs, and accelerators also means it can scale workloads from legacy systems to next-generation processors. This lets engineers update models independently of the processors, knowing there will always be a hardware platform to efficiently execute their AI.

AI, Predictive Maintenance, and the Road to ROI

One final requirement of predictive maintenance is that someone shepherd the living AI models. Ideally this is a member of the technical staff where the system is installed. But it’s unlikely that such a role exists at many companies today.

That makes tools such as Intel® DevCloud—which provides a sandbox for testing models before deploying them on operational hardware—a valuable platform. For organizations looking to support P&AM systems on their own, DevCloud can be used to experiment with analytics and diagnostics capabilities or learn how AI will integrate with legacy devices.

By combining development platforms like DevCloud and OpenVINO-optimized deployment hardware, automation companies can finally implement usable predictive maintenance systems. And in doing so, they can smooth the road to ROI.